

SPAWAR



*Systems Center
San Diego*

TECHNICAL DOCUMENT 3043
September 1998

**Data Fusion Project
Arrangement:
Final Report**

The United States/Australia
Memorandum of Agreement on
Radar Activities

W. C. Torrez, Editor
W. J. Yssel, Editor

Approved for public release;
distribution is unlimited.

19981112 013

TECHNICAL DOCUMENT 3043
September 1998

**Data Fusion Project
Arrangement:
Final Report**
The United States/Australia
Memorandum of Agreement on
Radar Activities

W. C. Torrez, Editor
W. J. Yssel, Editor

Approved for public release;
distribution is unlimited.



Space and Naval Warfare Systems Center
San Diego, CA 92152-5001

**Reproduced From
Best Available Copy**

SPACE AND NAVAL WARFARE SYSTEMS CENTER
San Diego, California 92152-5001

H. A. Williams, CAPT, USN
Commanding Officer

R. C. Kolb
Executive Director

ADMINISTRATIVE INFORMATION

The work detailed in this document was initiated under the auspices of the United States/Australia Memorandum of Agreement for Cooperation on Radar Activities. Project leaders were W. J. Yssel of the Space and Naval Warfare Systems Center, San Diego (SSC San Diego) and D. J. Percival of the Defence Science and Technology Organisation, Salisbury, Australia.

Released by
G. L. Wilham, Head
Test and Evaluation Branch

Under authority of
J. A. Audia, Head
Signal Exploitation and
Information Management
Division

Data Fusion Project Arrangement: Final Report

Project Leaders:

W. J. Yssel (United States)

D. J. Percival (Australia)

United States Participants:

W. J. Yssel	Space and Naval Warfare Systems Center, San Diego	(619) 553-2330
W. C. Torrez	Space and Naval Warfare Systems Center, San Diego	(619) 553-2020
N. Nguyen	Space and Naval Warfare Systems Center, San Diego	(619) 553-2340
S. W. Shaw	SRI International	(540) 859-4494

Australia Participants:

D. J. Percival	Defence Science and Technology Organisation (DSTO)/ Surveillance System Division	+61 88 259-7166
K. A. B. White	Defence Science and Technology Organisation (DSTO)/ Surveillance System Division	+61 88 259-7116
W. T. Johnson	Defence Science and Technology Organisation (DSTO)/ Surveillance System Division	+61 88 259-7141
I. W. Dall	Defence Science and Technology Organisation (DSTO)/ Surveillance System Division	+61 88 259-5419
A. J. Shellshear	formerly DSTO	
D. J. Kewley	formerly DSTO	

EXECUTIVE SUMMARY

This report describes the accomplishments of the Over-the-Horizon Radar (OTHR) Data Fusion (DF) Project Arrangement (PA), initiated under the auspices of the United States/Australia (US/AS) Memorandum of Agreement (MOA) for Cooperation on Radar Activities. This report also contains Project Arrangement background information and an explicit statement and discussion of each of the three primary Project Arrangement Objectives. The three primary objectives of the Project Arrangement are as follows:

1. Share OTH radar data collected and exchange relevant technical information.
2. Jointly formulate suitable data fusion approaches, developing algorithms capable of exploiting overlapping coverage.
3. Jointly devise and conduct experiments to test and evaluate data fusion algorithms.

The structure of the report is designed according to how each of these objectives was accomplished as a result of concerted Project Arrangement activity. Section 1 is an introduction.

Section 2 contains Project Arrangement background information and fulfills objective 1 by describing the procedures that were instrumental in the exchange of data and technical information between the United States and Australia for over-the-horizon radar systems.

Section 3 provides an overview of each country's over-the-horizon radar system resources and capabilities and sets the stage for the main activities of the Data Fusion Project Arrangement.

Section 4 addresses objective 2 concerning the formulation of track data fusion algorithms in the overlapping coverage areas of multiple over-the-horizon radars.

Section 5 addresses objective 3 and describes data collection experiments and the test and evaluation of track data fusion algorithms applied to the collected data. In particular, three track data fusion algorithms are described in detail in sections 4 and 5, (a) Multiple ROTH (Relocatable Over-the-Horizon Radar) Track Data Fusion, (b) DATAFUSE, and (c) ALICE (Always Logical in Connecting Events).

Section 6 addresses objective 2 for the case of associating and fusing multiple OTHR tracks with microwave radar tracks, a potentially important future operational requirement as defined by the user communities. US/AS Project Arrangement scientists are exploiting common approaches and data sets for the on-going implementation and Test and Evaluation issues of these more general data fusion algorithms, thereby fulfilling objectives 1 and 3 as well.

Finally, section 7 introduces methods for the performance assessment of general data fusion algorithms by using the notion of global data fusion. Dr. I. R. Goodman of the Space and Naval Warfare Systems Center, San Diego (SSC San Diego), has recently developed theoretical methods based on these ideas. Other US/AS researchers have applied these methods to OTHR networks incorporating additional surveillance assets. This approach will be useful in the future design of measures of effectiveness for the test and evaluation of the proposed track data fusion algorithms for the multisource-multisensor information sensor networks of which OTHRs play a major role.

The methods described in this report will enable a quantitative evaluation of the performance characteristics of quite general data fusion algorithms, that is, those that combine both classification and

positional information in a multitarget and multisensor environment. This approach introduces diverse measurements to form "best" estimates. Therefore, the work described in this report is oriented to the application, rather than the theoretical aspects, of the development and performance evaluation of data fusion algorithms. In particular, this work has been applied in the case of fusing OTHR tracks with OTHR tracks as well as fusing microwave radar tracks with tracks of targets of interest generated by OTHR systems.

Both United States and Australian systems have tremendous capabilities as independent sources of OTHR surveillance technology. Each system's interoperability and fusion architectures, if properly designed, give the users a surveillance capability of tremendous coverage and flexibility. Much is known about the mathematical properties of optimal design in the classical case of a single sensor and a single target. Similarly, in the case of multiple, nonhomogeneous sensors and a *single target*, many partial results are known. However, in today's information-rich, highly connected computing environments, there has been little, if any, investigation of the performance characteristics of combining information derived from a multitarget environment and multiple nonhomogeneous sensors. The approaches described in this report will play a key role in the design of optimal data fusion algorithms and will be pursued in future research and development programs for over-the-horizon radars.

CONTENTS

EXECUTIVE SUMMARY	iii
1. INTRODUCTION	1-1
1.1 SCOPE OF THE REPORT	1-1
1.2 INTEGRATED DATA FUSION SYSTEMS	1-1
1.3 CURRENT STATE OF OTHR DATA FUSION ALGORITHMS	1-2
1.4 STRUCTURE OF THE REPORT	1-4
1.5 SECTION REFERENCE	1-4
2. BACKGROUND	2-1
2.1 OBJECTIVES OF THE US/AS MOA DATA FUSION PROJECT ARRANGEMENT	2-1
2.2 EXCHANGE OF DATA FROM OTH RADARS AND EXCHANGE OF TECHNICAL INFORMATION	2-1
2.2.1 OTHR and Related Data Copied from the U.S. by AS/DSTO	2-1
2.2.2 Exchange Relevant Technical Information	2-2
2.2.3 AS/DSTO Living Notebook	2-4
2.3 JOINT FORMULATION OF TRACK FUSION ALGORITHMS	2-4
2.4 JOINT EXPERIMENTS FOR THE TEST AND EVALUATION OF DATA FUSION ALGORITHMS	2-5
2.5 SECTION REFERENCE	2-5
3. OPERATIONAL AND EXPERIMENTAL OTH RADAR RESOURCES	3-1
3.1 U.S. RELOCATABLE OVER-THE-HORIZON RADAR (ROTHR)	3-1
3.2 JINDALEE OPERATIONAL RADAR NETWORK (JORN)	3-5
3.3 SECTION REFERENCES	3-11
4. DATA FUSION IN OTH RADARS	4-1
4.1 MULTIPLE ROTHR TRACK DATA FUSION (MRTDF)	4-1
4.2 APPROACHES TO DETECTION, TRACKING, AND FUSION FOR MULTIPLE OTH RADARS	4-3
4.3 DATAFUSE AND ALICE IN AS SYSTEMS	4-56
4.4 A COMPARISON BETWEEN US AND AS DATA FUSION APPROACHES	4-78
4.5 SECTION REFERENCES	4-81
5. STATISTICAL SUMMARY OF ROTHR DATA FUSION	5-1
5.1 MRTDF ACCEPTANCE TESTING	5-1
5.1.1 MRTDF RESULTS	5-1
5.2 MEASURES OF EFFECTIVENESS FOR MRTDF	5-3
5.2.1 Introduction and Background	5-3
5.2.2 Probability of Multitrack	5-3
5.2.3 Position Accuracy	5-3
5.2.4 AOU and Containment	5-4
5.2.5 Course and Speed Accuracy	5-4

5.2.6	Track Maintenance	5-5
5.2.7	Summary and Conclusions	5-5
5.3	SECTION REFERENCES	5-5
6.	THEORY AND PRACTICE OF MULTISOURCE-MULTISENSOR DATA FUSION	6-1
6.1	RECENT DEVELOPMENTS IN FUSING MICROWAVE RADAR TRACKS WITH RELOCATABLE OVER-THE-HORIZON RADAR (ROTHR) TRACKS	6-1
6.2	TRACK FUSION AND COORDINATE REGISTRATION FOR MULTIPATH OVER-THE-HORIZON RADAR	6-10
6.3	MULTIPLE TARGET LOCALIZATION AND TRACK ASSOCIATION FOR OVER-THE-HORIZON RADAR	6-44
6.4	SECTION REFERENCES	6-63
7.	CONCLUSIONS AND A VISION OF THE FUTURE	7-1
7.1	GOODMAN-MAHLER APPROACH TO GLOBAL DATA FUSION	7-1
7.2	MATHEMATICAL FORMULATION OF GLOBAL DATA FUSION	7-1
7.3	SUMMARY AND CONCLUSIONS	7-2
7.4	SECTION REFERENCE	7-3
APPENDIX: LIST OF ACRONYMS		A-1

Figures

1.2.1.	Integrated Data Fusion System	1-2
3.1.1.	Multiple Relocatable Over-the-Horizon Radar (ROTHR) Track Data Fusion	3-1
3.1.2.	Overlapping Coverage of ROTHR Systems	3-2
3.1.3.	Propagation Characteristics of a Single ROTHR System	3-3
3.1.4.	Example of Range-Doppler Map for a Single ROTHR System	3-4
5.1.1.	P-3 Track for 31 May 1995 2230-2400Z	5-2
5.1.2.	CERAP Track for 1 June 1995 0001-0019Z	5-2
5.2.1.	A is the center of the ROTHR AOU. C is the true position of the target aircraft. Normalized Miss Distance is AC divided by AB	5-4
7.3.1.	Current State of the Mathematical Formulation of Optimal DFAS	7-3

1. INTRODUCTION

1.1 SCOPE OF THE REPORT

This report describes activities in the general area of Over-the-Horizon Radar (OTHR) Data Fusion, undertaken as part of the activities under the Data Fusion (DF) Project Arrangement (PA) of the United States/Australia (US/AS) Memorandum of Agreement (MOA) for Cooperation on Radar Activities.

1.2 INTEGRATED DATA FUSION SYSTEMS

What is data fusion and why is it important? Doug Kewley, Defence Science and Technology Organisation (DSTO), Australia, has stated, "Data fusion is the combination of data from multiple sources to provide enhanced information, quality, and availability over that which is available from any individual sensor." The U.S. Joint Directors of Laboratories (JDL) Data Fusion Group defines data fusion as "a continuous process dealing with the association, correlation, and combination of data and information from multiple sources to achieve refined entity, position, and identity estimates, and complete and timely assessments of resulting situations and threats, and their significance." Another illustrative definition is the "the integration of information from multiple sources to produce the most specific, comprehensive, and unified data about an entity."

What are the basic steps or building blocks of data fusion processes? Data fusion is composed of several overlapping levels. Each level contains several processing stages. They are known as alignment, association, correlation, attribute data fusion or classification, and analysis data fusion or situation assessment.

ALIGNMENT: Processing of sensor measurements to achieve a common time base and spatial reference.

ASSOCIATION: A process by which the "closeness" of sensor measurements is computed.

CORRELATION: A decision-making process that employs an association technique as a basis for allocating sensor measurements to the fixed or tracked location of an entity. In this regard, let us also define a **CORRELATION-TRACKER** as a process that generally employs both correlation and fusion component processes to transform sensor measurements into updated states and covariances arising from multiple hypotheses for entity tracks.

ATTRIBUTE DATA FUSION OR CLASSIFICATION: A process by which some level of identity of an entity is established, either as a member of a class, a type within a class, or a specific unit within a type.

ANALYSIS DATA FUSION OR SITUATION ASSESSMENT: A process by which the distributions of fixed and tracked entities are associated with environmental, doctrinal, and performance data.

Why perform data fusion of OTHR data? There are many benefits that can be derived from fusing OTHR data from multiple sources. If that fusion process is performed at the radar sites where detailed information and knowledge of the ionospheric conditions exist, improvement can be obtained in the following areas:

- a. **Improved Target Detection:** Returns from an OTH radar may propagate via different ionospheric paths that possess different and independent fading characteristics. Association and

fusion of data from these different propagation modes increases the timeliness of initial detection while eliminating multiple tracks from a single target.

- b. **Improved Track Maintenance:** Multiple, independent sensors will also have different ionospheric fading patterns and characteristics, and this will allow target tracks to be maintained through temporary losses by one of the member sensor suites. If these sensors have different aspects on the target, the chance of target loss due to insufficient Doppler is reduced.
- c. **Improved Track Accuracy:** Redundant data from independent sensor systems increase the accuracy by providing more independent "looks" at the target. Spatially separated OTHR systems allow for near orthogonal inputs that improve localization (increasing positional accuracy, while decreasing the Area of Uncertainty [AOU]) and improve course and speed estimates of the target.
- d. **Better Coordinate Registration:** With proper geometry, complementary sensor systems can aid in resolving ambiguities and uncertainties in the ionospheric mode identification process and help resolve the multiple hypotheses of associated systems. Systematic positional differences between tracks from the separate radars can be used to improve the estimation of ionospheric parameters.

Furthermore, even more striking improvements can be obtained if one admits data from auxiliary organic sources such as ground-based microwave radars and Global Positioning Satellite (GPS) or Federal Aviation Administration (FAA) and the international ADS (Automatic Dependent System) truth data sources. Figure 1.2.1 is a depiction of the ideal Integrated Data Fusion System, which would interrelate data fusion with other information processes such as Detection and Tracking and Coordinate Registration and which would include external data sources to provide a common operational picture in a time-critical manner.

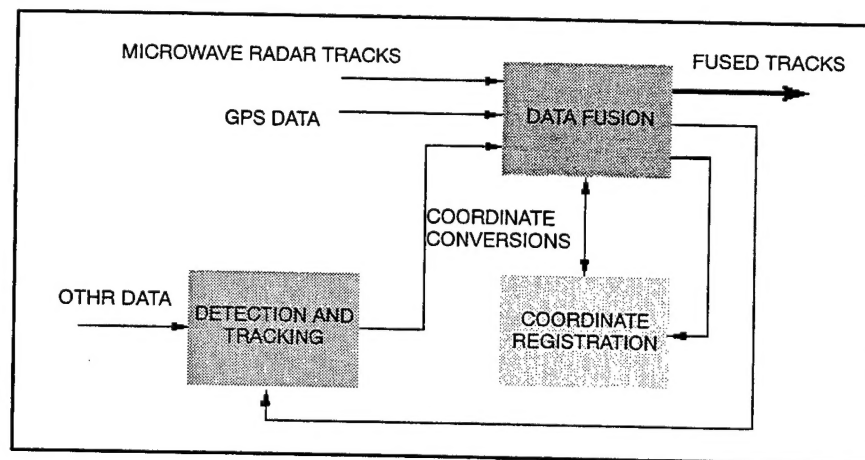


Figure 1.2.1. Integrated Data Fusion System.

1.3 CURRENT STATE OF OTHR DATA FUSION ALGORITHMS

A number of factors are relevant for understanding how both countries have implemented data fusion technologies within experimental and operational systems. Some of these factors relate to the nature of the respective national missions of U.S. and Australian OTH radars, and to the resultant demands and requirements of the user communities for the operation and effective use of system resources. Other factors relate to the chronological evolution of OTH radar technology research and development and to its application for the realization of mission objectives. However, in spite of

somewhat divergent missions, the activities pursued in accordance with the MOA have had tremendous impact on the development of both U.S. and Australian implementation of OTH radar.

The most dominant feature that both implementations have in common is that they are characterized primarily by their design as a distributed multisensor network, with OTH radars as a single, though important, entity. Because of the vast primary coverage areas (the entire north coast of Australia; for the U.S., 4,000,000 square miles in the Caribbean Basin and northern South America), joint commanders using these systems will want to use all organic assets, including OTHRs, ground-based millimeter or microwave radar networks, fixed-wing aircraft, aerostat vehicles, and other surveillance assets.

The most ideal system envisioned by the users is one that allows data from sources of widely varying physical characteristics to be fused, resulting in the Integrated Data Fusion System introduced in section 1.2. This system requires alignment of the sensor measurements in both time and space. In the case of fusing track data from two or more OTH radars, the problem can be done in a number of ways, for example, by aligning radar slant (or ground) track measurements so that they share a common time base and a common spatial reference. At an even more basic level, but possibly achieving more accuracy in the statistical estimates, the range-Doppler data could be fused. However, such an implementation would incur increased complexity in the alignment process of evaluating multiple hypotheses due to the ambiguous Doppler, the range-Doppler coupling, the ionospheric multipath, and the spatial separation of the radars. For the U.S. Relocatable Over-the-Horizon Radar (ROTHR) system, the approach aligning radar slant (or ground) track measurements so that they share a common time base and a common spatial reference was adopted, and this approach and its implementation will be described in detail, in section 4.1, and in sections 5.1 and 5.2. The approach fusing the range Doppler data, although not adopted, has been of keen interest and will be discussed in section 4.2.

In the case of fusing OTH radar tracks with microwave radar tracks in the areas of overlapping coverage, an alignment technique has been implemented for Australia's JFAS (Jindalee Facility—Alice Springs) radars. This approach was adopted by the Australians in their development of the DATAFUSE algorithm, which will be described in section 4.3. This technique requires that tracks, generated by the ground-based microwave radars, be transformed into radar slant coordinates (assuming a propagation path), and then the OTH track in slant coordinates is fused with the transformed coordinates of the ground-based radar track. The drawback of this technique is making good estimates of the covariances. These covariances arise from assumptions of the propagation paths used to register the coordinate transformations so as to achieve sensor alignment in coordinate space. The same alignment drawback exists if one, instead, transforms the radar slant coordinates from the OTH radars to ground coordinates, and then fuses those measurements directly with those from the ground-based microwave radars. This approach, described in section 6.1 (ground track fusion from two or more OTHRs with microwave radars), has been demonstrated by using commercial off-the-shelf software, such as the Advanced Tactical Workstation (ATW), developed by Raytheon/Texas Instruments. Some relevant comments on a comparison of MRTDF (track fusion of two OTH radars) and DATAFUSE will be discussed in section 4.4. Application of ATW for the fusion of ROTHR tracks with microwave radar tracks will be discussed in detail in section 6.1. Current research efforts pursued by the Australians focus on multipath track fusion for a single OTHR, and together with the nondeterministic likelihood calculus of S. Mori, will ultimately be applied for the track data fusion of OTHRs with microwave radar tracks in the regions of overlapping coverage. These efforts will be discussed in section 6.2.

1.4 STRUCTURE OF THE REPORT

Section 2 of this report contains project background information, including an explicit statement and discussion of each of the three primary US/AS MOA DF PA Objectives. Section 3 gives an overview of US/AS OTHR system resources and capabilities and sets the stage for the main activities of the Data Fusion Project Arrangement.

Three primary objectives of the US/AS MOA DF PA are as follows:

1. Share OTH radar data collected and exchange relevant technical information.
2. Jointly formulate suitable data fusion approaches, developing algorithms capable of exploiting overlapping coverage.
3. Jointly devise and conduct experiments to test and evaluate data fusion algorithms.

The structure of the report is based on a description of how each of the objectives was accomplished during the course of the DF PA activity. In terms of these objectives, sections 4 and 5 most directly address objective 2 for the formulation of track DFAs in the overlapping coverage areas of multiple OTHRs. In these sections, three DFAs—MRTDF, DATAFUSE, and ALICE—are described.

Fulfilling objectives 1 and 3 has been achieved as a byproduct of

- a. the formulation of these DFAs, and
- b. the design of experiments for the test and evaluation of these DFAs as they were and are being implemented into the operational systems.

Section 6 again addresses objective 2 for the formulation of track DFAs so as to associate and fuse track information from multiple OTHRs with microwave radar tracks, an important future operational requirement as defined by the user communities. Again, U.S. and Australian DF PA scientists are exploiting common approaches and data sets for the on-going implementation and T&E issues of these more general DFAs, thereby fulfilling objectives 1 and 3.

Finally, section 7 introduces methods for the performance assessment of general DFAs by using global data fusion, developed principally by I. R. Goodman and extended by other U.S. and Australian researchers. This approach will have use in the future design of measures of effectiveness for the T&E of the proposed track DFAs for the multisource-multisensor sensor networks mentioned in section 1.3.

1.5 SECTION REFERENCE

Yssel, W. 1994. "OTH Sensor Fusion," Naval Command, Control and Ocean Surveillance Center (NCCOSC) RDT&E Division (NRaD) White Paper, San Diego, CA (September).

2. BACKGROUND

2.1 OBJECTIVES OF THE US/AS MOA DATA FUSION PROJECT ARRANGEMENT

As stated previously, the primary objectives of the US/AS MOA Project Arrangement in Data Fusion are:

1. Share OTH radar data collected and exchange relevant technical information.
2. Jointly formulate suitable data fusion approaches, developing algorithms capable of exploiting overlapping coverage.
3. Jointly devise and conduct experiments to test and evaluate data fusion algorithms.

2.2 EXCHANGE OF DATA FROM OTH RADARS AND EXCHANGE OF TECHNICAL INFORMATION

2.2.1 OTHR and Related Data Copied from the U.S. by AS/DSTO

Australia's Jindalee Over-the-Horizon Radar Network (JORN) network will require fusion of tracks from multiple OTHRs. The U.S. is the only source of multiple OTHR data, and under the auspices of the MOA, a variety of data from the U.S. OTHRs and related sources has been made available to the Wide Area Surveillance Division (WASD), formerly the High Frequency Radar Division (HFRD). These data present an ideal opportunity to test track fusion algorithms on real overlapping OTHR data and will allow a comparison of Australian and U.S. approaches.

In August 1994, a dual OTH system trial was conducted. The U.S. Navy ROTHr in Virginia and the U.S. Air Force OTHB in Maine were operated with overlapping coverage of the Caribbean region. The FAA also provided data from an FAA microwave radar located in Puerto Rico.

The August 1994 data can be categorized as follows:

- Over-the-Horizon-Backscatter (OTHB) slant tracks and OTHB ground tracks
- ROTHr Virginia (VA) slant tracks and ROTHr (VA) ground tracks
- FAA Puerto Rico (PR) microwave radar tracks

The three types of data above are available for 3 August 1994 and 5 August 1994 only.

May 1995 Data

In May 1995, the two ROTHr systems (Texas and Virginia) were running as individual ROTHr systems. P-3 flights were conducted in support of the Acceptance Testing of the Texas site. ROTHr Virginia data were collected so that the Multiple ROTHr Track Data Fusion (MRTDF) system could be tested.

For the May 1995 set, the following data are available:

- ROTHr VA slant tracks
- ROTHr VA Coordinate Registration (CR) tables
- ROTHr TX slant tracks

- ROTHr TX CR tables
- MRTDF ground tracks and single ROTHr ground tracks
- CERAP (Panama) microwave radar tracks
- P-3 GPS ground positions

On 18 May 1995, a P-3 was flown south from Jamaica, overshooting Panama, and returning to Jamaica.

On 31 May 1995, a P-3 was flown in two approximately 10-sided overlapping polygons in a region between Jamaica and Panama. (Please note: For 31 May data, ground tracks are from the two ROTHrs operating as single systems, i.e., MRTDF was not used.)

The data sets described above have been stored on CD-ROM by AS/DSTO. Details on the nature and contents of these data sets are provided in the "living notebook," a world wide web (www)-based collection of HTML (hypertext markup language) files developed by AS/DSTO. This notebook will be described in section 2.2.3. Items described in the living notebook include hypertext links to the following information:

- File listing of U.S. data
- Slant track file format (part of Logical Record Identifier [LRID] 223)
- Slant track file format (all of LRID 223)
- Ground track file format (LRID 225)
- Coordinate Registration Tables (LRID 701)
- Explanation of CRT fields (LRID 701)
- CRT data for one Virginia Dwell Illumination Region (DIR) at one time instant
- Propagation mode IDs
- ROTHR transmit and receive site information

2.2.2 Exchange Relevant Technical Information

Numerous technical documents, reports, and conference proceedings papers have been exchanged between DF PA participants. Of special note is the active participation of PA principals and PA colleagues in presenting and publishing papers in the proceedings of various symposia and conferences and publishing papers in the open literature on various aspects of OTHR:

FUSION 98: First International Conference on Multisource-Multisensor Data Fusion

Recent Developments in Fusing Microwave Radar Tracks with Relocatable Over-the-Horizon Radar (ROTHR) Tracks. W. J. Yssel and W. C. Torrez

1997 IRIS National Symposium on Sensor and Data Fusion, MIT Lincoln Labs

Performance Characteristics of Data Fusion Methods with Application to Independent Surveillance Systems. W. C. Torrez and W. J. Yssel

Australian Data Fusion Symposium, Adelaide 1996

Hierarchical Clustering for OTHR Track Fusion. M. Mohandes, R. E. Bogner, and A. Bouzerdoun

Multi-Way Track Association for Over-the-Horizon Radar. J. Zhu, R. E. Bogner, and K. J. Pope

Robust Estimation of Ionospheric Density using Multisensor Sounding Data. E. Skafidas, D. J. Percival, and R. J. Evans

Measures of Effectiveness for Multiple ROTH Track Data Fusion (MRTDF).

W. J. Yssel, W. C. Torrez, and R. A. Lematta

AIMTEC 1994

Automatic Multi-Sensor Data Fusion Processing using Fuzzy Logic. B. J. Jarvis and D. J. Kewley

Digital Imaging Computing Techniques and Applications—DICTA 1995

Multiradar Track Fusion using the Hough Transform. M. Mohandes, R. E. Bogner, A. Bouzerdoun, and K. J. Pope

IEEE Transactions on Aerospace and Electronic Systems

An Association Metric for Track Association of Multimode OTHR Tracks, M. L. Soutcott and R. E. Bogner.

Multiple Track Association for Over-The-Horizon Radar. R. E. Bogner, A. Bouzerdoun, K. J. Pope, and J. Zhu

IEEE Signals, Systems and Computers Conference 1993

Evaluation of a Model Based Data Fusion Algorithm with Multi-mode OTHR Data. I. W. Dall and A. J. Shellshear

Multiple Track Association for Over-the-Horizon Radar. R. E. Bogner, A. Bouzerdoun, K. J. Pope, and J. Zhu

International Conference on Neural Networks

A Comparison of Neural Networks and Statistical Methods for Track Association in Over-the-Horizon Radar, 1995. J. Zhu, R. E. Bogner, A. Bouzerdoun, K. J. Pope, and M. L. Soutcott

A Neural Network Approach to Multiradar Track Fusion, 1996. M. Mohandes, R. E. Bogner, A. Bouzerdoun, and K. J. Pope

International Conference on Radar

Track Association in the Presence of Multi-mode Propagation. UK, 1992. I. W. Dall and D. J. Kewley

Performance of an Improved Model Based Data Fusion Algorithm for OTHR Data. I. W. Dall and A. J. Shellshear

Signal Processing, SPIE, 1994, vol. 2233

Application of Neural Networks to Track Association in Over-the-Horizon Radar. J. Zhu, R. E. Bogner, A. Bouzerdoun, and M. L. Soutcott

TTCP Subgroup K

Association of Tracks from Multiple OTHR. I. W. Dall

9th National Symposium on Sensor Fusion, Monterey, CA. 1996

Multiple Relocatable Over-the-Horizon Radar (ROTHR) Track Data Fusion (MRTDF).
W. J. Yssel, W. C. Torrez, and R.A. Lematta

DSTO/AFOSR Signal Processing Workshop, Victor Harbor, June 1997

Multipath Track Fusion for Over-the-Horizon Radar. D. J. Percival and K. A. B. White

SPIE Signal and Data Processing of Small Targets, San Diego, July 1997

Multipath Coordinate Registration and Track Fusion for Over-the-Horizon Radar.
D. J. Percival and K. A. B. White

Miscellaneous

Implementation Issues for Datafuse at JFAS. Andrew Shellshear and Ian Dall

ALICE: A Fuzzy Logic Approach to Automatic Multi-Sensor Data Fusion. W. T. Johnson
and D. J. Kewley

A complete bibliography of relevant OTHR literature is given in the AS/DSTO living notebook previously mentioned. The living notebook and its contents are described in the next section.

2.2.3 AS/DSTO Living Notebook

The "living notebook" is a collection of www pages, designed to quickly document DF PA activity as well as allow for the easy electronic access and exchange of relevant technical information between U.S. and AS DF PA participants. Pages include PA documents, presentations, descriptions of track fusion algorithms, and experimental data sets. In addition, the notebook contains links to the US/AS MOA and an extensive bibliography of relevant DF PA OTHR books, technical reports, and conference proceedings. The notebook will continue to play an important role in providing valuable reference material and descriptions of work in progress in an easily accessible web-page format.

2.3 JOINT FORMULATION OF TRACK FUSION ALGORITHMS

Due in large part to the imminent installation of the second ROTHR system in 1995 and the U.S. user requirement that the operators be provided with a single "fused ROTHR" track, the U.S. pushed forward to fuse radar tracks from multiple ROTHR sites in radar slant coordinates. The details describing the approach, application, measures of effectiveness (MOEs) and acceptance testing of the multiple ROTHR track data fusion (MRTDF) algorithm are given in sections 4.1, 5.1, and 5.2.

SRI International, funded through Rome Laboratories, investigated the fusing of OTHR data at the range-Doppler map level. The results of this investigation are summarized in section 4.2. A report describing these results will be generated and distributed in the near future.

DSTO researchers from the Wide Area Surveillance Division (formerly HFRD) took multiple slant tracks from a ROTHR P-3 target and used DATAFUSE to associate and fuse the tracks and compared the results with the ROTHR output. A second DFA, called ALICE (Always Logical in Connecting Events) was also developed to allow fuzzy logic principles to be applied to track fusion. The results and application of these DFAs are described in section 4.3.

Also under the PA, a study was initiated to evaluate the improvement in coordinate registration by the use of track fusion generated information in the region of overlapping coverage areas. This work was initially undertaken by Tom Berkey, a Utah State University faculty member, who was a visiting researcher at Rome Labs during 1995–1997. This work will be continued and extended by SSC San Diego resident scientists.

2.4 JOINT EXPERIMENTS FOR THE TEST AND EVALUATION OF DATA FUSION ALGORITHMS

In the case of fusing track information for multiple OTHRs, DSTO researchers from WASD provided a list of MOE metrics to the U.S. for the evaluation of DFAs, such as MRTDF. SSC San Diego, formerly the Naval Command, Control and Ocean Surveillance Center (NCCOSC) RDT&E Division (NRaD), incorporated these metrics into the test and evaluation of Multiple ROTHr Track Data Fusion (MRTDF) in numerous experiments as part of the Texas ROTHr acceptance testing. The description of the OTHR data sets analyzed, and the results of the test and evaluation of the MRTDF approach, have been extensively shared with AS scientists.

Discussions between WASD and SSC San Diego during September 1996 culminated in a philosophy of how to fuse OTH radar target tracks with microwave radar tracks on the same target (in the overlapping coverage areas). The OTHRs possess large coverage areas and long holding times, but OTHR target tracks possess large AOUs. The microwave radars, of course, have shorter holding times, but very accurate localization capabilities. Section 6.1 describes U.S. efforts to apply DFAs (implemented in the Advanced Tactical Workstation [ATW], which provides such capabilities), while section 6.2 describes independent Australian efforts that also offer the capability to associate and fuse OTHR target tracks with microwave radar tracks. By coincidence, both of these efforts apply theoretical approaches to DFAs described by S. Mori in a number of papers (*cf.* bibliography provided by DSTO in the “living notebook”). The work of Mori, an application of a nondeterministic likelihood calculus technique, appears to unify the U.S. and AS data fusion approaches, and future experiments based on the same OTHR data sets will most likely bear this out.

2.5 SECTION REFERENCE

Percival, D. J. et al. 1997. “Living Notebook World Wide Web Home Pages,” DSTO Salisbury, AS, (September).

3. OPERATIONAL AND EXPERIMENTAL OTH RADAR RESOURCES

3.1 U.S. RELOCATABLE OVER-THE-HORIZON RADAR (ROTHR)

ROTHR is a long-range, bi-static, high-frequency radar surveillance system tasked with tracking air targets over large land and ocean areas. The long ranges, together with the relatively low altitudes of the targets, require that the radar look beyond the line-of-sight, in other words, over-the-horizon (OTH). This is accomplished by refracting the signal off the ionosphere to points beyond the horizon. The radar transmitter is designed to radiate 200 kilowatts of RF power from 16 log-periodic, amplitude tapered, phased-array antennas. The ROTHR system surveys a coverage area consisting of azimuths and ranges up to 80 degrees and 2000 nautical miles. The U.S. Navy operates two ROTHR systems (AN/TPS-71 ROTHRs) located in the United States. The first one, located in Virginia, looks south over the Gulf of Mexico and the Caribbean. The second one, which has recently been installed in south Texas, looks to the southeast and provides coverage of the Gulf and Caribbean. The geometry of the two radars and the region of overlapping coverage are shown in figures 3.1.1 and 3.1.2.




Figure 3.1.1. Multiple Relocatable Over-the-Horizon Radar (ROTHR) Track Data Fusion.

The primary mission of the ROTHR system is to provide continuous detection and monitoring of suspected air smuggling activities in the Caribbean and Eastern Pacific Regions. Track data fusion is also an issue of common interest with Australian OTH radar development interests, so a secondary mission of the program is to interact with Australian Defence engineers and scientists in the research and development of engineering and operational enhancements of OTH radar.

These OTH radar provide three features that complement the capabilities of the existing surveillance assets for detecting and tracking aircraft operating in the Gulf and the Caribbean. First, OTH radar have relatively large coverage areas; second, with a suitable choice of additional OTH radar locations (e.g., Puerto Rico), they can be made to look at airspace deeper into South America (figure 3.1.2). This is the region where drug traffic originates. Third, in contrast to the look-up capability of conventional ground radar, OTH radar provides a look-down capability. This feature enables the OTH radar to detect drug-carrying aircraft that follow flight paths through valleys in mountainous regions. Such flight paths would be masked from ground-based microwave radar. The look-down capability of OTH radar provides a better detection capability.




Figure 3.1.2. Overlapping Coverage of ROTHR Systems.

ROTHR is an over-the-horizon radar system that looks beyond the horizon by bouncing its signal off the ionosphere to illuminate targets and by having the reflected signal return in a similar manner. Amateur radio operators will appreciate that long ranges are possible through the use of this concept. ROTHR operates between 5 to 28 MHz depending on the time of day, distance to the target, and the current state of the 11-year sunspot cycle. The ground also provides a very strong return signal (much stronger than that from the target), which requires Doppler processing to distinguish the target from the ground clutter.

Targets must have some discernible radial velocity component toward or away from the fixed radar installations. (It doesn't have to be very much). The ionosphere is not a copper plate. It is multilayered, dynamic, and nonstationary in time and space. This gives rise to multiple radar returns from a single aircraft target. These returns arrive from different radar ranges and slightly different radar azimuths due to differing elevation angles (see figure 3.1.3). The ROTHR system monitors the ionosphere with sounders and maintains a current ionosphere model (CIM) and generates coordinate registration tables (CRTs) for each expected combination of modes in order to transform the radar




Figure 3.1.3. Propagation Characteristics of a Single ROTH System.

from radar coordinates (radar range and radar azimuth) into ground coordinates of latitude and longitude.

The main technical challenge arises from the fact that the OTH radar relies on the ionosphere to propagate the radar signal to and from the target. For each detected target, ROTH measures the slant range, the slant azimuth, and the Doppler of the target. Target tracks are typically required in ground coordinates (e.g., latitude and longitude). Transformation of the slant measurements to ground coordinates requires knowledge of the path followed by the radar signal. Such a path is also referred to as the propagation mode, and it specifies the reflecting ionospheric layers and the number of bounces the radar signal undergoes on its way to the target and back to the ROTH receiver. Several such propagation modes are feasible, and for each mode there is a corresponding transformation of the slant measurements to ground coordinates. It is difficult to predict the path followed by the radar signal because the ionospheric conditions depend on many variables (e.g., the frequency of radar operation, the geographical area of operation, the time of the day, the sunspot activity, etc.), and the conditions continually vary with time. If the mode is not selected correctly, then the ground track computed from the slant track will be inaccurate. The errors in the ground tracks due to incorrect mode selection will thus give rise to registration errors, and then it will be difficult to correlate the ground tracks of one OTH radar with another.

Figure 3.1.4 shows the amplitude (in dB) of a typical signal return from a single beam as a function of range and range rate of the target. The middle of the Doppler axis on this map represents zero Hz, and the large hump around this value represents stationary returns backscattered from the earth's surface; this is called ground clutter. On this map, one can clearly see the target with a well-resolved peak with positive range rate, but it is spread in range. This spread in range with a single ROTH results because the signal return is made up of returns coming back via different ionospheric paths (as discussed in the previous paragraph).

The formation and maintenance of tracks are done by associating or combining returns from a target on a given ionospheric mode in radar coordinates, but the association or correlation of slant tracks originating from the same target aircraft is done in ground coordinates. Target track association is performed by (1) bringing each set of possible associations to ground by each of the available ionospheric modes, and (2) looking for clustering of the ground track positions. The best selection of mode assignments is judged by a chi-square test in latitude, longitude, course, and speed. Target




Figure 3.1.4. Example of Range-Doppler Map for a Single ROTH System.

track correlation is performed by taking a weighted average of the individual slant track positions to form a ground track position.

With the addition of the second ROTH site in Texas, with a coverage area that significantly overlaps that of the Virginia radar, the users of the data requested that they be given one ROTH track per target, not one track from the Texas ROTH and another track from the Virginia ROTH. The straightforward approach would have been to build a system that would associate and correlate the ground tracks from the two ROTH sites and provide a composite ROTH track. But we already had a single ROTH fusion system running, and it seemed prudent to expand it to include a second and third ROTH. The thought was that since ROTH is a Doppler radar and targets are lost when they fly perpendicular to the look direction of the radar, what could be better than to have inputs from a near orthogonal second ROTH to maintain coverage on the target. Also, since course and speed errors always seem to increase as the target's orientation approaches a broadside aspect to the radar, this second data input would improve course and speed accuracy.

Errors in target position are often dominated by errors in range due to the misidentification of the ionospheric mode. A second ROTH looking near perpendicular to the first to help correctly identify the ionospheric mode being used, thus provides significant gain for detection and tracking, as will be shown in section 4.

3.2 JINDALEE OPERATIONAL RADAR NETWORK (JORN)

JORN

Jindalee Operational Radar Network

Presented by Cheryl Zakaria

The Essential Features of the JORN Project


<u>The Two Radars</u>	<u>The Frequency Management System</u>
<u>The JORN Co-Ordination Centre in South Australia</u>	<u>A Network of Communication Links</u>

The Two Radars

Radar 1 is located near Longreach in the state of Queensland. It has a 90 degree field of view (along a boresight of 325T) extending from the Solomon Islands across to the North-West Cape.

Radar 2 is located near Laverton in the state of Western Australia. It has a 180 degree field of view (along a boresight of 350T) stretching roughly from Cairns out to the middle of the Indian Ocean, approximately 12.5 million square kilometres. Here there is an extra arm at both the transmit and receiver sites.

The scan area overlaps to ensure that no target can utilize tangential tracking to avoid detection.




(JFAS - Jindalee Facility at Alice Springs is the operational OTH radar that is situated in the Northern Territory & is the upgrade of the pilot radar Stage B Jindalee. JFAS is not part of JORN)

Each radar consists of two sites, a transmitter site and a receiver site, which are located approximately 100 kilometres apart.

The Transmitter Site:

The transmit antennas are vertical log-periodic curtains in a uniform line array on a wire ground mat, which extends about 20 m. in front of the array. The installed length of the transmit arrays is 0.4 km. (At Laverton there are two arrays at 90 degrees to each other.)

The 28 high power amplifiers at each site have a 20 KW CW rating and are capable of transmitting a maximum of 560 KW. The amplifiers can be coupled into either a 28 element highband array or a 14 element lowband array. In the 14 element lowband array the amplifiers are coupled pair-wise, through combiners to deliver 40 KW into each antenna. The arrays can be split into half arrays and electronically beamsteered through 90 degrees. The transmitted waveform will be linear sawtooth frequency modulated carrier wave.



Transmitter Site at Laverton W.A.

The Receiver Site:

The antenna array is 3 km in length. It consists of 480 pairs of monopole antennas, with each pair connected to its own front end receiver. The wire ground mat extends about 50 m forward and assists in the concentration of returning signals.



Receiver Site at Laverton in W.A.

The strict time synchronization between the transmit and receive sites is performed by the use of Atomic standards with periodic reference back to global positioning systems data.

GEC Marconi is responsible for building the radar transmitters and receivers as well as the Stage 1 & 2 radar control hardware and software, the data control distribution, the digital waveform generators, and the transmit drives and receiver controllers.

The Frequency Management System :

It is the frequency management system that determines which frequencies will yield the best signal-to-noise ratio in the band of frequencies.

The system comprises of:-

Spectrum monitor and background noise monitor which will be used to ascertain clear channels

and background noise respectively.

Background scatter sounders, located at the radar sites, which are needed to measure ionospheric propagation characteristic.

Low powered mini-radar, like a mini-JORN, and a passive channel evaluator which will provide the capability to independently access the proposed operating frequencies more directly.

A series of oblique incidence sounders, vertical incidence sounders and transponders are strategically located around the northern coastline. They provide an analysis of ionospheric structure and assist in automatically establishing the ground truth of targets.

The frequency management system as well as the vast antenna arrays which are the dominant physical feature of the JORN radar sites at Laverton & Longreach are being built by Radio Frequency Systems Pty. Ltd.

The JORN Co-Ordination Centre:

The JORN Co-Ordination Centre located at Edinburgh RAAF base will be the hub of the network. It consists of an operations centre, software support and training facility and an administration and security facility. The predicted manning for the JORN Co-Ordination Centre for 24 hour operations is 145 service personnel.



The JORN Co-Ordination Centre

The actual building is 3 800 square metres. The heart of the Co-Ordination Centre is the central operations room where operators will use at least 20 Digital Corporation Alpha workstations. The signals will be fed from the Co-Ordination Centre's mainframe "backroom" where seven massive capacity Digital Alpha APX 7000's capable of processing in excess of 500 megaflops of information being relayed from the JORN remote radar sites.

All sensor control including the frequency management, all of the target tracking and much of the raw signal processing will be carried out centrally at this Co-Ordination Centre.

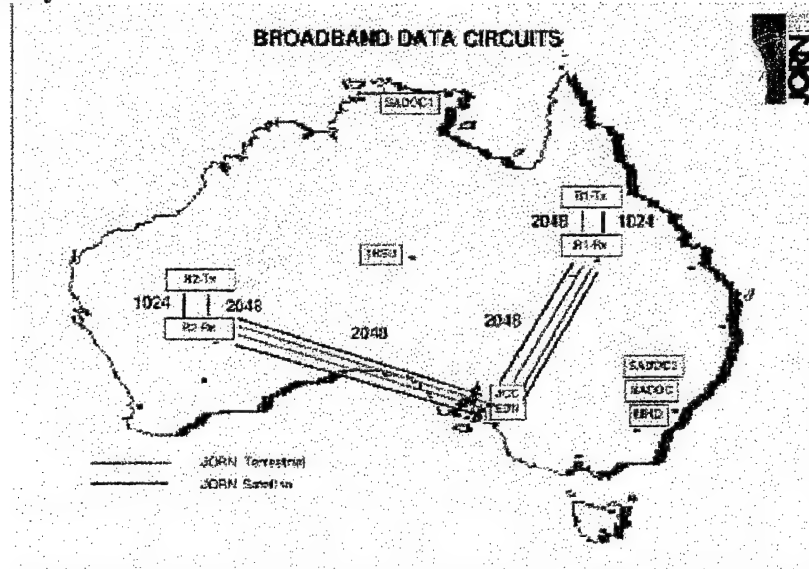
The software needed for this system to operate is overwhelming. Approximately 900 000 lines of ADA code is being developed. Telstar Systems Pty. Ltd., a company set up in 1990 as a 60:40 joint venture between Telecom Australia and Lockheed Missiles & Space Company is responsible for 90% of this software.

In the event of a catastrophic failure occurring at the JORN Co-Ordination Centre there are autonomous or secondary operations centres located at each of the radar receiver sites.

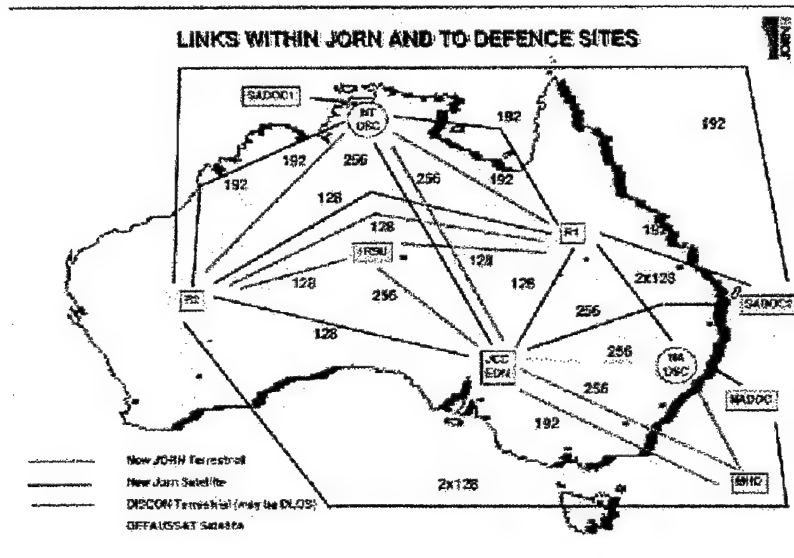
The Network of Communication Links:

Communications are vital to the operation of each of the JORN radars and the operation of the network. Local area communications are required for each radar to ensure co-ordination and timing between the transmitters and receivers. The primary links are provided by fibre optical

Wide area communications is being designed around the need to pass 4 megabytes/second bandwidth of data between the radar sites and the Co-Ordination Centre on 2MB links. The primary medium for transmission of this data will be TELSTRA hired terrestrial megalinks with the JORNSATCOM providing back-up. During non-contingency times the satellite transponder will be on cold standby.



Finally this wide area communications network includes links from the JORN Co-Ordination Centre to its external agencies, the organizations that are being provided information, such as the bureau of meteorology, customs & immigration agency. These links will be primarily be DDS Fastway point-to-point terrestrial links with some priority links backed up by cold satellite transponder links.



3.3 SECTION REFERENCES

Yssel, W., W. Torrez, and R. Lematta. 1996. "Multiple Relocatable Over-the-Horizon Radar (ROTHR) Track Data Fusion (MRTDF)," *Proc. 9th Nat. Symp. Sensor Fusion*, pp. 37-45, Monterey, CA, (October).

Zakaria, C. 1997. "JORN World Wide Web Home Pages," (January).

4. DATA FUSION IN OTH RADARS

4.1 MULTIPLE ROTHTR TRACK DATA FUSION (MRTDF)

There are several points in the ROTHTR processing stream where data fusion principles can be applied.

- a. ROTHTR Ground Track Fusion: One could simply take the output ground track information from two ROTHTR systems and combine them at the ground position level. This Ground Track Fusion could select the "best" of the tracks by computing a weighted average of the separate tracks or by picking the strongest or most consistent track.
- b. ROTHTR Slant Track Fusion: The ROTHTR already does a form of data fusion in its handling of multiple mode slant tracks. This Slant Track Fusion is done by allowing the slant tracks that are associated with a given target to be brought to ground using all possible combinations of existing ionospheric modes. The set of modes that provides the "best fit" (most consistent picture) is assigned to the set of tracks and a weighted average of the target's parameters is used to represent the target state. Fusion of tracks from multiple ROTHTRs can be accomplished in a similar fashion by expanding the space of consideration to include slant tracks from all radars that can be associated with the same candidate target and by using all of the available modes from each of the radars.
- c. ROTHTR Peak to Existing Slant Track Fusion: A third approach could involve the fusing of peaks to existing slant tracks. Each time the radar visits a Dwell Illumination Region (DIR), the set of peaks detected during that dwell are reviewed to see which peaks can be associated with existing slant tracks. After the existing slant tracks have been checked, any leftover peaks that exceed a signal-to-noise (SNR) threshold are used to begin new slant tracks. Peaks that are below the threshold are discarded. At this point, a new level of association could be initiated to see if (1) any of the lower SNR peaks represent a detection on an ionospheric mode that has not yet received sufficient energy to initiate a slant track or (2) any of these peaks represent detections of an existing target on a second radar. In either case, these low SNR peaks could be brought into the track family or be the first step in initiating new slant tracks to be associated with the track family. An additional level of complexity could be pursued at the Range-Doppler Map level. Since an approximation of the state of the target is known and the ionospheric modes have been identified, one could search the Range-Doppler maps to see if low-level peaks exist that could represent detections on other modes or other radars and that would not have been strong enough to have passed the peak threshold.
- d. ROTHTR Peak-to-Peak Data Fusion: There is an extension to the process just described above. For any peaks that were left after attempting to associate them with existing tracks, one could take the unassigned peaks and keep them in a database pool and use them to attempt to initiate possible tracks in the common area of multiple radars, operating on the available ionospheric modes.
- e. ROTHTR Base Level Range-Doppler Map Data Fusion: There seems to be a maxim in data fusion that states, "The closer to the sensor that the data fusion takes place, the more accurate the output product." However, there is no "free lunch." The price for such an implementation is in increased complexity in the alignment process (multiple hypotheses due to the ambiguous Doppler, the Range-Doppler coupling, the multiple ionospheric paths, and spatially separated radars). Also complicating the picture are the large bandwidth and extensive data storage

requirements needed to handle all of the Range-Doppler maps from both radars for the time interval necessary to initiate tracks.

The scheduled Initial Operational Capability (IOC) of the ROTHr-Texas installation in the latter part of FY 95 called for swift action to meet the users' desire to have ROTHr report out single tracks, representing and containing contributions from all ROTHr sensors. Realizing that the intent of data fusion is to provide a product that shows a measurable improvement over its constituent parts, it seems prudent to take a cautious approach. This required building on successes, measuring the incremental improvements as one moves from the simple to the complex, and taking the time to discover the weaknesses and investigate the pitfalls before stumbling onto them. Data fusion is a fairly recent developmental area, and certainly its application to multiple OTH radars represented an exciting new challenge.

Algorithms for accomplishing ground track data fusion from two or more ROTHrs are easily available. There are probably more than a dozen systems that could handle the task, and one of them, the Advanced Tactical Workstation (ATW) is discussed in section 6.1 for the fusion of OTHr tracks with microwave radar tracks. With the bringing together of the Operational Control Centers (OCCs) for the three proposed ROTHr systems (Virginia and Texas already online and Puerto Rico, proposed IOC 1999), into a common area, Slant Track Data Fusion became the chosen approach for the first step in ROTHr data fusion. Slant tracks from overlapping DIRs, whether they be from the same or different radars, will be merged in the same way that slant tracks in overlapping DIRs, or multi-mode tracks from the same DIR, are merged in the current ROTHr system. The vehicle for merging slant tracks is the mode identification and linking process. Newly verified tracks, as they are brought to ground, are tested against all existing track families. If the verified track can be associated with an existing track family, it is correlated with that family and is assigned the mode that is most consistent with the existing track families. If the verified track can be associated with an existing track family, it is correlated with that family and is assigned the mode that is most consistent with the existing track. Sometimes this will entail the reassignment of the ionospheric modes among all of the family members. If the newly verified track cannot be associated with an existing track family, it is assigned the strongest ionospheric mode and begins a new track family. The MRTDF IOC was completed shortly after the Texas ROTHr system came online, and the MRTDF/Texas Acceptance Test Results are discussed in section 5.1.

4.2 APPROACHES TO DETECTION, TRACKING, AND FUSION FOR MULTIPLE OTH RADARS

EXECUTIVE SUMMARY

The objective of this program was to develop improved methods for fusing data from multiple over-the-horizon (OTH) radars. This program was motivated both by need and opportunity. The need for improved tracking and fusion technology arises from new operational requirements for tracking smaller, slower, and more maneuverable targets, as well as a desire to place radars in locations with non-ideal clutter conditions. The opportunity to realize these new operational requirements comes about in part because of recent deployments of OTH radar sites with overlapping coverage in the United States and Australia. By synergistically applying simultaneous data from two independent radars, far greater tracking performance should be possible. In addition, the continuing growth of computing power enables us to pursue much more advanced, albeit computationally complex, techniques for tracking and fusion.

Within this setting, SRI undertook a research and development program to study and experimentally implement a new generation of OTH sensor fusion algorithms. The key result of our work has been the development of an innovative design that unifies target detection, tracking, and fusion within a single procedure. This approach has the potential both to reduce labor demands by fully automating track association between radars, and to enable the radar to detect and track smaller, slower, and more agile (maneuvering) targets.

The table below summarizes the benefits of our multi-sensor fusion approach.

Benefit	Enabling feature of multi-sensor fusion approach
<ul style="list-style-type: none">• Small target sensitivity• Tracking in clutter	By combining the complete (unthresholded) data set from multiple radars, tracking of smaller targets in clutter is possible.
<ul style="list-style-type: none">• Maneuvering target tracking	By tracking directly in the target's natural state-space, target maneuvers are more continuous and can be followed more easily.
<ul style="list-style-type: none">• Track accuracy	Continuous target location information from two radars yields tighter error ellipses.
<ul style="list-style-type: none">• Automation	Each target produces a single track, independent of the number of viewing radars. The ionospheric model parameters are calculated as a byproduct of tracking.

The key technical results of this project were:

- Development of an efficient method for mapping the entire sensor data set to target state space. An important outcome of this work is a collection of algorithms for efficiently transforming three-dimensional (3-D) radar sensor data into the natural 4-D coordinate system of targets.
- Establishment of computational feasibility. We estimate that a single 300-MHz Intel P-II class computer would be sufficient for real-time operations of the low-level fusion technique presented here.
- Demonstration of tracking through substantial target maneuvers, including continuous

tracking of targets through 180-degree turns.

- Demonstration of significant sensitivity gains. Data from the P3 experiment indicates SNR gains of 5 to 10 dB, depending on the accuracy of the ionospheric models and resulting coordinate registration.
- Determination that the algorithms achieve SNR gain even with relatively significant errors in ionospheric registration. Our statistical analysis indicates that up to 30 km registration errors by the ionospheric modeling process can be tolerated with less than 3 dB of sensitivity loss.
- Determination that inter-sensor target, clutter, and noise are statistically uncorrelated.
- Determination that ionospheric modeling errors are statistically uncorrelated.

1 INTRODUCTION

1.1 Background

The fusion systems reflect the view that tracking and multi-radar fusion are independent processes. There are several practical reasons for choosing the two-stage approach, including:

- *Reduced data requirement:* The first-stage tracking process eliminates the vast quantity of data from consideration in fusion. This data reduction has been necessary to stay within the computational limitations of affordable computers.
- *Mode assignment not required during tracking:* Associating radar measurements between two OTH radars requires that both measurements be located in ground coordinates. This, in turn, requires that a propagation mode be assigned to each measurement. Current OTH systems defer mode assignment to the latest possible stage so that tracks can be formed independent of any assumed ionospheric path.
- *Historic reasons:*
 - (a) Single-radar tracking algorithms are already deployed and well understood. It has been more practical to incrementally augment these systems by adding a second, largely independent fusion layer on top of existing tracking modules.
 - (b) *Availability of slant-to-ground conversion tools.* The current practice of combining multi-mode tracks of the same target during the slant-to-ground conversion closely resembles the multi-radar fusion problem. Each mode of propagation for a single target can be thought of as an individual sensor. The mode-linking procedure in place at ROTHr exhibits all of the characteristics of a fusion algorithm as discussed below, including statistical combination of evidence, generation and testing of multiple model hypotheses, and numerous heuristics intended to reduce computation in the search procedure. This procedure could easily be applied to the problem of fusing multiple radars as well as multiple modes.

The vast improvements in commercial computing have greatly changed the landscape of assumptions under which OTH tracking and fusion systems were designed, however. Since the Navy's Relocatable OTH Radar (ROTHR) was deployed in the late 1980s, computing capacity has grown by a factor of at least several hundred¹. By applying the additional available computing power to the fusion problem, much more aggressive and complete signal processing approaches can be pursued. Specific to the fusion process, the processing power enables more data to be processed, more model hypotheses to be generated and evaluated, and the ionospheric models and mode assignments to be made at the earliest stage of processing. These factors have led us to consider fusion during the tracking stage rather than as a secondary process. To show how these concepts are merged, the following sections review some of the current practice in both tracking and fusion technology.

¹ The Digital Equipment Corporation VAX 780 was a 1-MIP computer that cost approximately \$250,000. Today's four-processor 333 MHz P-II computer costs \$25,000 and produces the equivalent of 1000 VAX MIPs, achieving an overall price/performance improvement of about 10,000:1.

1.1.1 Current Tracking Technology

Traditional approaches to target tracking rely on a two-stage process whereby raw data are first thresholded by means of a detection process. This detection process is then followed by a data association stage that links individual detections over time. Examples of this approach include the Probabilistic Data Association Filter (PDAF) [Bar-Shalom, 1974], Joint PDAF [Bar-Shalom and Tse, 1975], and variations on the Multiple Hypothesis Tracking (MHT) algorithm [Reid, 1979; Blackman, 1986]. Use of these approaches has been largely motivated by the substantial computational savings afforded by their hierarchical down-selection of data. By reducing the quantity of information that need be considered at each stage in the process (i.e., detection, association, filtering), fewer alternatives need be considered in any search procedure. The disadvantage of this approach is the loss of information, and in particular the potential for missing dim targets that do not regularly reach the required detection threshold. Maneuvering targets are also difficult to track. To reduce false alarms, standard tracking approaches require a target to demonstrate consistency in location and speed before a target declaration is made. Maneuvering results in apparently inconsistent target returns, and if the maneuvers are at all significant, a track is never declared.

More aggressive sensor-level processing approaches to tracking have been developed to improve target sensitivity. Motivated by requirements for small-target detection, and enabled by rapidly improving processor technology, several new approaches based on dynamic programming were developed for several application areas [Wishner et al., 1981; Barniv, 1985; Arnold, 1986; Barniv and Kella, 1987]. These techniques perform target detection and track formation from within a single optimization procedure, and thereby improve target/noise discrimination performance if target motion models are available. Various performance studies on these techniques indicate approximately 3-5 dB of improved target sensitivity [Barniv and Kella, 1987].

1.1.2 Current Fusion Technology

1.1.2.1 Defining Fusion

The phrase "multi-sensor fusion" raises confusion from the merger of two separate but non-exclusive goals that are set for most fusion systems. The first of these goals can be stated as:

Produce a fused measurement that is more accurate and/or more reliable than any single measurement taken alone.

Some fusion systems fail to achieve this goal in pursuit of one that is more basic, yet often just as difficult to achieve. That goal is:

From a group of sensors, produce, at most, one report for any single target.

Another way to state this dual definition is that fusion consists of both a data association problem and an estimation problem. Although these goals are not mutually exclusive, one or the other typically takes precedence. In an effort to manage large amounts of sensor data and present a unified view to the operator, the method of combining sensor measurements may be overlooked. Conversely, in attempting to optimally estimate target state from a group of measurements, the track maintenance and association problem may not receive adequate

attention. By combining the tracking and fusion problems into a single process, we will devote equal attention to the problems of *data association* and *optimal estimation*.

The manner in which fusion is accomplished depends on the compatibility between sensors. We can categorize fusion algorithms as low-level, mid-level, or high-level. However, it should be noted that a continuum of possibilities exists, and no single algorithm belongs entirely in one category or another.

Low-level fusion: When the sensors to be fused use the same sensing modality and are precisely registered, such as a multi-band infrared imaging sensor, the fusion process can often be accomplished at the raw sensor data level. In these applications the sensor measurements are already aligned, and hence the association process is straightforward. The technical challenge in this case is to develop methods that exploit multivariate statistical models for the sensor suite.

Mid-level fusion: When the sensors measure different physical quantities but take their measurements in the same format, such as an infrared imaging sensor and a visible light camera, mid-level fusion takes place. Such a process may involve remapping sensor data onto a common measurement space or extracting key features from the data.

High-level fusion: This approach is often appropriate when both the physical modality and the format of the sensors differ. Such an example might be the fusion of an infrared imaging sensor with a direction-finding acoustic array. Not only do the sensors measure very different physical quantities, there is no clear mapping from the sensors onto a common measurement space. In these situations, preliminary interpretations must be made independently from each sensor and symbolically combined through the use of additional domain knowledge.

From this definition of fusion algorithms we can anticipate the most natural form for fusing data from multiple OTH radars. OTH radars utilize the same sensing modality and collect those measurements in identical formats. However, because of the variation in sensor acuity with range, azimuth, and heading of the target, widely separated OTH sensors will not, in general, produce measurements that are directly comparable. The resulting fusion algorithm will lie somewhere between the low and middle levels, indicating that the sensors may be fused at the raw data level with the intervention of some mapping onto a common measurement grid. High-level fusion, involving abstract features that are far removed from the original data, is not, at least in principle, required.

1.1.2.2 Algorithm Characteristics

Regardless of sensor type or the specific algorithm, all fusion systems have certain features in common. These features may be present in varying degrees, but to accomplish the goals of association and estimation, any algorithm must involve some statistical estimation, deal with multiple hypotheses in a nondeterministic manner, and carry out some symbolic inference.

Uncertainty: The problem of estimating a parameter from one or more error-prone measurements has a long history of study. In fusion systems, the process may be formulated within the bounds of traditional Bayesian estimation, or it may make use of recent developments in evidential reasoning or fuzzy set theory. Regardless of the specific implementation, a fundamental fusion algorithm feature is the ability to quantitatively assess uncertainty in individual measurements and to combine those uncertainties into an overall value of target belief.

Nondeterminism: Fusing disparate sensor measurements requires accurate models of both the environment and the sensor. These models often involve parameters that are unknown or unobservable. The result of this incomplete knowledge of the world is a system that postulates solutions under various assumptions and tests the veracity of those assumptions once the answer is derived. This may be viewed as nondeterminism in the sense of a nondeterministic automaton. There are two means by which this nondeterminism may be expressed in a fusion system:

- The system iteratively postulates solutions and tests the results. If unacceptable results are generated, the system enters the generate-and-test cycle again. This cycle is repeated until an acceptable solution is found or time runs out.
- The system generates a number of potential solutions and sorts through the resulting answers. More than one competing solution may be maintained until a report is made.

We prefer the latter approach because of operational requirements for real-time operations. There is no guarantee that the iterative approach will terminate with an acceptable solution, and there is no way to predict the processing time required to generate a correct solution. With the multiple-hypothesis approach, a constant processing load is maintained, regardless of problem difficulty or uncertainty.

Inference: Because the information required to assess a potential fusion solution may be other than numerical, many fusion systems incorporate an element of logical inference in addition to statistical computation. This inference may be executed within a formal, rule-based expert system, or it may be manifested as a loosely coupled set of heuristics used to limit the size of a solution space or to focus on a particular approach to the problem. We have found it useful to formalize the inference procedure in order to manage application of heuristics. Without such organization, the fusion system can become unwieldy, inefficient, and difficult to debug.

1.2 Problems Unique to OTH Fusion

In addition to the general fusion issues discussed above, OTH radar presents some unique problems to the sensor fusion algorithm designer. These problems are largely the result of the ionospheric propagation path, which can alter a target's response drastically from one sensor to another. In addition, the wide geographic disparities between receiver sites in a typical multiple OTH configuration complicate the data association problem. For these reasons, existing, off-the-shelf fusion solutions are not readily adaptable to the OTH problem. Some of the specific problems that may be encountered are:

- *Spread clutter:* One of the advantages of tracking with two or more radars is the ability to resolve actual target velocity given two estimates of radial velocity. In an OTH radar, ionospheric clutter at low Doppler frequencies may entirely mask returns from a target whose motion is substantially tangential. This may result in detection of a target in one sensor and not in another, even though the target's location is within the joint coverage area.
- *Disparate ionospheric models:* Because of differing propagation paths, returns from a single target observed in two different radars may undergo differential distortions in terms of amplitude, range offset, and apparent azimuth. Any fusion algorithm must

model the statistics of these variations and take them into account when scoring an association. For example, a target may fade in one sensor and remain visible in another, and the fusion algorithm must allow for this possibility.

- *Multimode propagation*: Depending on transmitter power and ionospheric conditions, reflections from a single target may arrive along several propagation paths. Each of these paths produces a separate target response in the data. The results must then be combined into a single track. In a single-sensor tracker, each mode may be tracked separately and fused at the post-tracking stage. However, across two or more sensors, the potential exists for fusing two different modes, resulting in ghost associations.

Although these factors can be seen as obstacles, they are also incentives for implementing fusion. The same mechanisms are present in single-sensor systems, and the resulting detrimental effects can be lessened by considering multiple OTH sensors instead of one. Although the possibility of a missed detection complicates the fusion process, it is only through fusion that an intermittent target can be continuously tracked across ionospheric fades and through spread clutter. The result is that OTH radars are at their most effective if considered jointly rather than as single, independent assets.

1.3 Other Recent Work in OTH Fusion

Several efforts have been carried out to develop OTH track-level fusion algorithms. In the United States, joint development by NRaD and Raytheon led to a slant-to-ground fusion algorithm using modifications to the current ROTHr mode-linking algorithm [Yssel, 1994]. This was an appropriate near-term solution because it allowed fusion to be implemented with minimum modification to the current ROTHr software and operational procedures, and led to an immediate fusion capability between the ROTHr Virginia and Texas facilities.

Alphatech also performed experiments in fusion between Over-the-Horizon Backscatter (OTH-B) and ROTHr Virginia data under Rome Laboratory support [Kurrien, 1994]. Their experiences point out the problems in fusion of track reports from radar systems that were intended as stand-alone sensors. Not enough information is supplied in these tracks to form the sensor and propagation models' required hypotheses. In addition, if a track quality measurement is supplied at all, it cannot be directly related to a statistical model for target likelihood.

Extensive OTH track fusion investigations have been under way in Australia for some time. This work is characterized by the iterative, hypothesize-and-test approach [Kewley, 1994]. In addition, the Australian work attempts to exploit collateral information such as ground-based radar detections or intelligence reports in fusion process. To successfully incorporate this additional knowledge, researchers are focusing on techniques such as fuzzy sets for managing uncertainty and neural networks for recognizing patterns. Use of supporting information allows the fusion system to derive, in addition to simple estimates of target position and heading, the possible intent or nature of the target.

2 OBJECTIVES

Our goal has been to develop a seamless tracking and fusion architecture that provides performance benefits to both multi-site and conventional single-site OTH facilities. The resulting approach extends current OTH tracking technology in several significant ways. Specific innovations include:


- *Multi-sensor target detection and tracking directly on the sensor data.* To make maximum use of the sensor information, target tracks are formed on the raw, unthresholded sensor data. In addition to improving system sensitivity to weak targets, this approach improves track maintenance performance, especially against maneuvering targets.²
- *Tracking and fusion as a single, unified process.* The target detection and track formation processes are carried out within a single optimization procedure. This procedure takes into account the sensor data within the context of models for target cross-section, target motion, radar operating parameters, and ionospheric stability.
- *A fusion framework that invites other sensor types.* By first projecting the sensor data into a common target coordinate system, the proposed approach allows data from other sensors, including other classes of sensors, to be incorporated into the target detection and tracking process.
- *Registration information incorporated in the track formation process.* Coordinate registration information from terrain mapping, CREDO, or other techniques is used to assist the data association process, as well as to improve target position estimates.
- *Robust statistical methods that mitigate the effects of clutter.* To address the spatial and temporal nonstationarity of OTH data, special robust statistical analysis techniques are employed. These approaches identify regions (in space, Doppler, and time) of homogeneous statistical behavior, and develop appropriate non-Gaussian noise models. These statistical descriptions are then used to transform the measured target amplitudes into target likelihood measures.
- *Parallel processing exploited from the start.* Modern commercial parallel processors provide an opportunity for more aggressive exploitation of sensor data, provided that the algorithms are designed with parallelism in mind.

The low-level sensor fusion approach developed here involves three key steps: (1) projection of the data from all sensor sources into a common target coordinate system; (2) combination of the projected sensor information into probabilistic measures for target occupancy; and (3) formation and detection of tracks in target coordinate space.

² The point regarding maneuvering targets is not immediately obvious, and will be explained at greater length.

2.1 Sensor to Target-Space Projection

The operational model is illustrated in Figure 2-1. Two OTH radars provide independent observations of a partially overlapping region. The data stream at each sensor site is processed to resolve in slant range, azimuth, and radial velocity, and an amplitude estimate is formed for each resolution cell within this 3-D space. A representative example data set from an OTH facility is shown in Figure 2-2. In the example of Figure 2-2, data from three radar beams and three range resolution cells are shown over a period of several minutes. Note that to display inherently 4-D data (3-D plus time) on 2-D paper, the data must be doubly nested. In our chosen format, data from each radar beam is displayed within a broad horizontal stripe, with data from each radar observation period contributing a single thin horizontal line in this stripe. Range cell centers are indicated as bright vertical bars showing the strong zero-Doppler ground return signal. On either side of each bright center bar are the resolved Doppler cells for the corresponding range cell. Targets are shown as consistently brighter vertical trails over time at some non-zero radial velocity.



056-1

FIGURE 2-1 OPERATIONAL CONCEPT




FIGURE 2-2 EXAMPLE OTH GRAYSCALE DATA IN SENSOR COORDINATES

The unfortunate problem with the natural radar coordinate system is that it does not support fusion with other sensor sources, including other radars. In fact, combining data over time from a single radar is difficult if the radar operating parameters are varied. Radar coordinates are also an inappropriate model for target behavior, and are suboptimal for the track formation and track maintenance processes. This latter point is a particularly critical issue for maneuvering targets, where relatively modest changes in target state parameters (*e.g.*, target course) can result in dramatic and hence unmaintainable changes in the observed target Doppler.

To address these issues, our approach carries out all track processing in the natural 4-D coordinate system of targets: position (latitude, longitude), course, and speed. The 3-D sensor data is mapped into the 4-D space through a geometric transformation. Because of our goal of achieving full sensor data-level fusion, all sensor data is utilized; no pre-thresholding or other pre-selection is carried out. In our approach, a quantized grid in 4-D target coordinates is generated to cover the complete geographic coverage region of the radar, and the plausible range of target velocities.

The transformation function is given as the inverse of the familiar radar equations for predicting target range and Doppler. As illustrated in Figure 2-3, this inverse function maps each sensor resolution cell (range, Doppler, azimuth) into a curve in target coordinates (lat , lon , v_{lat} , v_{lon}).³ Note that this transformation results in a curve and not a point because the 4-D target state is not uniquely determined from (any single) 3-D sensor observation. The shape of the curve is given in 4-D by a constant position parameter, with speed and course parameters related by

$$\text{target speed} = (\text{radial velocity}) / \cos (\text{target course}) .$$

Although the target location is approximately known (within the accuracy bounds of OTH), the previous equation indicates that there is an ambiguity between target speed and course. As a result, all target state-space cells along the target course-speed ambiguity curve are equally supported by the sensor resolution cell, and hence each must be assigned an equal incremental observation scoring.

³ For clarity, a 2-D into 3-D mapping is illustrated, although the OTH application requires a 3-D to 4-D mapping (largely impossible to illustrate).




FIGURE 2-3 SENSOR TO TARGET COORDINATE MAPPING

Sensor resolution cells from all sensors project into similar curves in target space, where the shape of each curve is determined by the sensor viewing angle on the shared field of surveillance. Taken in combination, every *pairing* of sensor resolution cells from two radars results in a single point of intersection in target state-space (Figure 2-4); that is, a target's course, speed, and 2-D position are uniquely identifiable from a pair of 3-D observations.⁴ In fact, if the target signal-to-noise ratios were very high, and the target densities low (such that target associations across sensors are wholly unambiguous), the fusion process would be trivial from a single pair of observations. These optimistic conditions are not the reality of operational OTH radars, as their coarse position resolutions lead to high effective target densities, and their operational need to detect targets at low signal-to-noise ratios requires that many low-probability candidate data associations be investigated. More advanced processing stages, as described below, will be required to hypothesize and evaluate a large set of candidate target-data associations over some period of time.

⁴ Given range-Doppler ambiguities, this pairing can result in a small collection of intersecting points, as shown in the figure.




FIGURE 2-4 TWO SENSORS MAPPED INTO A SINGLE TARGET SPACE, RESULTING IN AMBIGUOUS ASSOCIATIONS

2.2 Estimating Target Occupancy Probabilities

The values projected into the target state-space resolution cells are transformed into likelihood ratios for target occupancy, defined as [Sharf, 1991]:

$$LR(a) = P(a|H_1) / P(a|H_0) ,$$

where $P(a|H_1)$ and $P(a|H_0)$ represent the probability for the observed sensor amplitude a under the assumptions of target presence and absence, respectively. A novel component to this work is the inclusion of observations from multiple sensors in this formulation; *i.e.*, the observation is now generalized to a vector form, and hence $P(a|H_1)$ and $P(a|H_0)$ are multivariate distributions. The statistical models necessary to carry this out in a multiple-sensor application do not currently exist, and must be developed under this effort. However, we have found evidence that the following general statements hold:

- The clutter will be statistically independent when observed by two OTH radars with approximately orthogonal viewing angles.
- The terrestrial noise sources (interference, etc.) may be strongly correlated between sensors.
- Target fades (*e.g.*, due to multipath interference) will be independently observed at each sensor.
- The effective radar cross-section of a target may be correlated between sensors, and may even be predictable, though perhaps in an impractical, complicated way.

Prior work in HF noise and clutter modeling indicates that non-Gaussian noise/clutter models are required, such as the mixed Rayleigh/log-normal distribution,

$$P(z) = (1 - e) P_r(z) + e P_{ln}(z) ,$$

where P_r and P_{ln} are Rayleigh and log-normal models, respectively, and e is the probability of a sample's having been contaminated from the (generally much stronger) log-normal distribution.

2.3 Target Track Formation and Detection

The ambiguous associations shown in Figure 2-4 are resolved through a tracking process illustrated in Figure 2-5. The approach to multi-sensor tracking formulates target detection as a search procedure that optimizes a target scoring function based on a target model in 4-D space. The scoring function is based on a target model that incorporates amplitude, as viewed through the likelihood ratio transformation given above, along with consistency in motion. Each target-space entry models target motion behavior through the update equation,

$$q(k+1) = A(k) q(k) + u(k) \quad (1)$$

where $q = [x, y, x', y']$ is the 4-D target state, A is the deterministic state update matrix,

$$A = \begin{bmatrix} 1 & 0 & dt & 0 \\ 0 & 1 & 0 & dt \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

and $u = [0, 0, x', y']$ is a random vector that introduces target maneuvers. Higher-order models (e.g., six-parameter models that include third-derivative ("jerk") components) are possible, although at a computational price.




FIGURE 2-5 CANDIDATE STATE TRANSITIONS BETWEEN ADJACENT INTERVALS

At each observation interval, all of the candidate states q evolve according to Eq. (1). Associated with each candidate target state is a cumulative likelihood value that provides an indication of confidence for target occupancy. To be strictly optimal, the algorithm must select

the maximum cumulative likelihood value over all hypotheses as the only target. However, in practice, the algorithm seeks out local maxima and compares them to a threshold, reporting those that exceed it. Tracks reported need not be consistent with each other, but may instead represent competing hypotheses as to target behavior. Further processing, such as a multiple-hypothesis tracking algorithm, can be used to select the best track from the reported hypothesis set, which is greatly reduced from the original hypothesis space. It is important to note that at each temporal stage of the tracking and detection algorithm, the hypothesis space is fully populated. All possible target states are evaluated, and hence there is no loss of optimality in reporting a track. A track report is merely the best estimate of target states at the latest time and may be revised in the future if conflicting data arrives.

The computation of cumulative likelihood requires a small search for each hypothesis. The bulk of processing time can be computed as the time to perform one of these searches multiplied by the total number of hypotheses. Thus, although there may be a large number of hypotheses, the computation time grows only linearly with this parameter. The key to reducing time complexity of the algorithm is to keep the search space for each hypothesis small. This can be done by choosing the correct dimensions for hypothesis space (*i.e.*, course and speed rather than v_{lat} v_{lon}), and by preprocessing the data to eliminate as much uncertain variation in registration and modeling as possible from one time step to the next.

3 TECHNICAL DISCUSSION

3.1 Theoretical Derivation

The dynamic programming algorithm (DPA) for tracking is a key step in our fusion approach. The DPA algorithm has been used successfully to detect and track small targets in single-site OTH radar applications, and its ability to enhance the detectability of small targets and to exploit target motion models also makes it ideal for the fusion problem. However, to make the best use of DPA tracking, we must first map the data from both sensors into a common reference frame. In this case, the common reference frame will be the space of all possible target hypotheses in ground coordinates. By a target hypothesis, we mean both the target's position on the ground and its instantaneous velocity in both the latitude and longitude directions. To most effectively utilize the DPA tracker, this mapping must also be done in a statistically meaningful way. In this section, we derive the probabilistic mapping procedure that allows us to apply our existing DPA code to the fusion problem.

The relevant statistic from the DPA perspective is the likelihood of target occupancy at any one hypothesis cell. In the following discussion, we represent a target hypothesis at time k by the symbol θ_k , the null target hypothesis as H_0 , and a single sensor measurement (in a single range-Doppler cell) at time k as z_k . In this notation, the DPA tracking procedure can be summarized as:

$$\arg \max_{\theta} \left\{ \frac{P(\theta_k | z_k)}{P(H_0 | z_k)} = f(p(z_k | \theta_k), p(z_k | H_0), \text{transition probabilities}) \right\}$$

In other words, we recursively choose the state θ_k at each time to maximize the likelihood ratio conditioned on the sensor measurements. Using Bayes' rule, this likelihood can be computed as a function of conditional probabilities as expressed on the right hand side of the above equation. For a more complete derivation and exposition on the DPA procedure applied to OTH radar, the reader is referred to [Shaw et al., 1996]

The single-sensor case can be expanded to multiple sensors by simply making the measurement variable, z_k , into a vector of two measurements, *i.e.*,

$$p(z_k | \theta_k) = p(z_k^1, z_k^2 | \theta_k)$$

where z_k^1 and z_k^2 are the measurements in sensors 1 and 2, respectively. Computing this joint conditional probability density function can be somewhat difficult, unless we can make the assumption that the measurements in sensors 1 and 2 are independent. We would like to answer the question

$$p(z_k^1, z_k^2 | \theta_k) \stackrel{?}{=} p(z_k^1 | \theta_k) p(z_k^2 | \theta_k)$$

If this assumption is indeed true, it will greatly simplify our computations. Section 3.4 details a correlation analysis of position errors between two sites in actual radar data. We have found that, for the periods of interest, there is no significant correlation between sites.

Usually we know, or can model the probability of making a particular sensor measurement conditioned on a target hypothesis. This is a fairly straightforward task and is the focus of considerable effort on the part of ionospheric scientists and radar designers. However, to apply DPA tracking in the space of target hypotheses, *i.e.*, on the ground, we must be able to compute the probability density function (pdf):

$$p(z_k^{\text{target}} | \theta_k)$$

This is the probability of making a certain measurement in ground coordinates conditioned on a target hypothesis, but we are not accustomed to computing this value directly. Instead, we have direct knowledge of the following quantities:

$$p(z_k^{\text{target}} | z_{ki}^{\text{sensor}} \theta_k), p(z_k^{\text{target}} | z_{ki}^{\text{sensor}} H_0)$$

which is the pdf of a measurement in ground coordinates given a single measurement in sensor coordinates and a target hypothesis or the null target hypothesis, and

$$p(z_{ki}^{\text{sensor}} | \theta_k)$$

which is the pdf of a sensor measurement at time k and measurement cell i , conditioned on a target hypothesis at time k .

So we must answer the following question: Can we compute the likelihood in ground coordinates given the information at hand?

First we must make the following assumption:

$$p(z_k^{\text{target}} | Z_k^{\text{sensor}} \theta_k) = 0,$$

where Z_k^{sensor} is the set of all possible sensor measurements $z_{ki}^{\text{sensor}}, i = 1, \dots, N$,

The interpretation of this assumption is that a particular target-space measurement can only result from a unique set of sensor measurements. In other words, the target measurement is unique. From now on, we will make this assumption although it is not absolutely true. Making this assumption simplifies the computation of a target-space measurement based on the entire collection of sensor measurements. If this assumption were not made, we would have to compute the probability of all other possible measurement space configurations each time we mapped a value from sensor to target space. This could become a combinatorial nightmare and would not contribute significantly to the final product, particularly since the result would most likely be a constant offset to the entire target-space hypothesis set.

Once the uniqueness assumption is made, the target space conditional pdf can be expressed as:

$$p(z_k^{\text{t}} | \theta_k) = p(z_k^{\text{t}} | Z_k^{\text{s}} \theta_k) p(Z_k^{\text{s}} | \theta_k)$$

This expression is the “sensor-to-target-space mapping.” It takes the set of sensor measurements, passes it through a conditional probability that we must model, and produces an equivalent set of target-space measurements. The remaining problem is to express the pdf of a target space measurement, z_k^t , conditioned on the entire data set in sensor space, Z_k^s , in a form that can be easily computed. To assist in this, we will make a final assumption, that the set of sensor measurements taken at one time is uncorrelated:

$$p(z_k^t | \theta_k, Z_k^s) = \prod_{i=1}^N p(z_k^t | \theta_k, z_{ki}^s), \quad p(Z_k^s | \theta_k) = \prod_{i=1}^N p(z_{ki}^s | \theta_k)$$

If this is true, then we can express the sensor-to-target-space mapping as:

$$p(z_k^t | \theta_k) = \prod_{i=1}^N p(z_k^t | \theta_k, z_{ki}^s) p(z_{ki}^s | \theta_k)$$

This is the form that we use to actually compute the sensor-to-target-space mappings. Expressed in log form, the expression is

$$\log[p(z_k^t | \theta_k)] = \sum_{i=1}^N \left\{ \log[p(z_k^t | \theta_k, z_{ki}^s)] + \log[p(z_{ki}^s | \theta_k)] \right\}$$

Thus, for each target hypothesis cell we take a group of measurement-space cells, discount each of them by the log of the conditional pdf $p(z_k^t | \theta_k, z_{ki}^s)$, and sum them together to give us the target-space value. This process is illustrated in Figure 3-1. The process is very much like a convolution in that we are passing a filter $p(z_k^t | \theta_k, z_{ki}^s)$ over the input data ($p(z_{ki}^s | \theta_k)$) to produce the output data, $p(z_k^t | \theta_k)$. The result is a mixture density in the target hypothesis space.

The critical unknown quantity in this formulation is the conditional probability $p(z_k^{target} | z_{ki}^{sensor}, \theta_k)$. This function incorporates both the positional uncertainty of a target on the ground that may have given rise to a particular sensor measurement, and the contribution to the “amplitude” of a given target hypothesis resulting from a particular sensor measurement. The shape of this function should relate to the statistical distribution of position errors for a given radar site. This information can be empirically determined. We will discuss this issue further in the statistical modeling section (Section 3.4) of this report.





FIGURE 3-1 ILLUSTRATION OF SENSOR-TO-TARGET-SPACE MAPPING

3.2 Fusion Processing Examples

This section describes, in practical terms, how the theoretical fusion process is carried out in the computer. We also present some examples to demonstrate the process. The first thing to note is that the geographical target hypothesis space is four-dimensional, consisting of latitude, longitude, and the latitudinal and longitudinal velocities. However, the sensor takes its measurements in a three-dimensional space consisting of range, azimuth, and range-rate. Thus, when a sensor measurement maps to the target hypothesis space, a single range-Doppler cell contributes to many ground hypothesis cells. The number of ground cells is determined by the spreading function $p(z_k^{target} | z_{ki}^{sensor} \theta_k)$, as described in the previous section.

We usually visualize this 4-dimensional space as series of disks, representing possible target bearings and velocities, embedded in a rectangular grid of geographic locations. Figure 3-2 illustrates this concept. We actually encode each 2-dimensional velocity-space slice from this data set as a rectangular grid rather than concentric circles, because the tracking software we have adapted for this project expects rectangular velocity components. However, the tracking algorithm might be more effective tracking in speed and bearing velocity components rather than in the rectangular coordinates imposed by latitudinal and longitudinal velocities.

What happens when a range-Doppler cell gets mapped into the 4-D ground coordinate space? Figure 3-3 illustrates this process. The missing dimension in the sensor measurement space is the azimuthal velocity. A single measurement of range rate constrains the hypothetical motion, *i.e.*, the total 2-D velocity vector, that a target could have. In rectangular, latitude/longitude velocities, the component in the direction of radar steer must be constant, whatever the target motion. Since the target velocity vector must project to a known value in this direction, the locus of hypothetical target velocities for a given Doppler measurement is a line in the direction normal to the radar steer. Thus, a single radar measurement-set maps to a set of lines in the 4-dimensional target hypothesis space.



058-2

FIGURE 3-2 ILLUSTRATION OF THE 4-DIMENSIONAL TARGET HYPOTHESIS SPACE




FIGURE 3-3 RADIAL VELOCITY MEASUREMENT OF A TARGET CONSTRAINS THE LOCUS OF POSSIBLE VELOCITIES

Figure 3-4 illustrates the mapping concept applied to two ROTHr dwells, one from Virginia, and one from Texas that contain the same known target. Although the hypothesis-space is 4-D,

we can only conveniently visualize two dimensions at a time. Therefore we usually display either a geographic space or a velocity space. The top row of images in Figure 3-4 shows all geographical bins sampled at the velocity of the known target. Note that the colored portion of the picture has the wedge shape characteristic of an OTH radar DIR. The bottom row of images shows the complete velocity hypothesis space sampled at the latitude and longitude where we know the target to be located. The significant values in these images form a disk because we have not attempted to compute the likelihood of absolute target speeds that we know are physically impossible, or at least highly unlikely. Note that each of the velocity space images on the bottom row consists of a set of stripes of varying intensity. Each of these stripes corresponds approximately to a Doppler cell in the original data. The dark stripe in the middle of each circle represents the blind speeds around zero Doppler that have been masked out. However, because of the spreading function that accounts for uncertainty in geolocation, there are actually many range cells making up each geographic bin in the hypothesis space. Thus, we see the smearing of several Doppler cells into each stripe. Figure 3-5 illustrates the velocity space diagrams in more detail. The last column in Figure 3-4 shows the geographic and velocity space slices of a hypothesis space that includes both the Virginia and Texas data. The target is, by far, the brightest feature in these plots. Note that the intersection of two moderately bright lines in velocity space has produced a single spot where the actual target velocity is known to be.

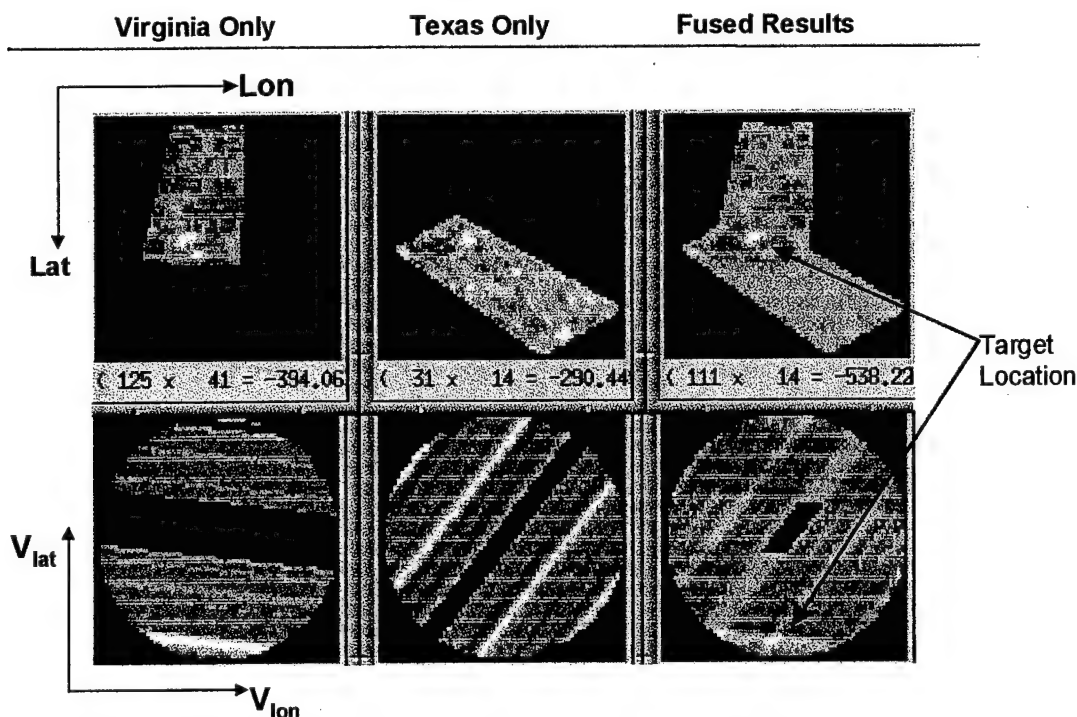


FIGURE 3-4 AN EXAMPLE OF FUSING DWELLS FROM TWO DIFFERENT OTH RADARS INTO THE SAME HYPOTHESIS SPACE

Another feature worth noting in Figure 3-4 is the “speckled” appearance of the bottom-row, velocity-space plots. These images were produced early in our research before the mapping software had been perfected. The speckled appearance is due to aliasing in the sensor-to-target-

space mapping. Because we are sampling one set of discrete measurements and converting them to values in a different lattice of discrete measurements, the potential for spatial aliasing exists. In any situation of this type, we must ensure that the bandwidth of the first data set is limited to the Nyquist range of the second data set. In later mapping software, we added a module to account specifically for this problem. We deliberately oversample the data, smooth it, then downsample onto the final resolution grid.




FIGURE 3-5 A DETAILED ILLUSTRATION OF THE TARGET HYPOTHESIS VELOCITY PLANE

Figure 3-6 illustrates the need for tracking as an integral part of the fusion process. This figure shows an example velocity hypothesis plane taken from a single geographic location at a single time. There is one known target in this resolution cell, but there are at least three separate peaks present in the data. Our mapping process actually creates more false peaks than would be present in single-sensor data because a new peak is generated for each pair of returns. If we were to threshold this data, as a standard constant false-alarm rate (CFAR) detection process would, two false peaks would emerge in addition to the actual target location. In the example of Figure 3-6, there are two target lines in the Texas data; these are actually range-Doppler ambiguous versions of the same target. Thus, our fusion technique helps to resolve range-Doppler ambiguities as well as improving target detectability. If we had only this single pair of dwells to on which to base a detection decision, all three peaks would probably be reported. However if we observe the data over time, the true target location becomes apparent.




FIGURE 3-6 A SINGLE TARGET VELOCITY PLANE SHOWING
MULTIPLE POTENTIAL TARGET PEAKS

Figure 3-7 shows a sequence of velocity planes taken over time. We have shown potential paths that a track might take through this data sequence. Although there are several possible targets in each frame, there is only one obvious target path that is consistent throughout the sequence; this identifies the true target position. The DPA tracking procedure excels at making this judgment automatically in a qualitative and efficient manner.

Figure 3-8 shows the results of processing a sequence of radar dwells that have been mapped into the ground-coordinate hypothesis space. This figure shows the cumulative log-likelihood values for three frames taken from a sequence of 27. It was not necessary to wait 27 frames for a detection; in fact, the target was clearly visible after the first few. However, we have shown a longer sequence to demonstrate the effect of DPA tracking on the data set. Note that in the first frame there are several candidate peaks, although the actual target peak is the brightest in the picture. As we move through time, the target peak grows in prominence and the others diminish. The bloom surrounding the target peak is characteristic of the DPA approach and can be detrimental to tracking closely spaced targets.

Notice that the shape of the target peak in Figure 3-8 is distinctly elongated in one direction. This is due to the difference in range and azimuth uncertainty and the relative revisit rates of the two radars. In this data set, the Texas radar was operating in staring mode over the area of interest. The Texas revisit rate was about 15 seconds, whereas the Virginia revisit rate was the standard 30-45 seconds. Because the Texas radar produced at least twice as much data, the error ellipse of the resulting target location estimate is narrow along the Texas radial direction and wide along the azimuthal direction. The range position of an OTH radar target is much less random than the azimuthal position, so our resulting cumulative score reflects this evidence.




FIGURE 3-7 SEQUENCE OF VELOCITY PLANES OVER TIME SHOWING THE CONSISTENT PATH DUE TO AN ACTUAL TARGET




FIGURE 3-8 CUMULATIVE LIKELIHOODS RESULTING FROM DPA TRACKING OVER TWO OTH RADARS MAPPED INTO THE FUSED TARGET HYPOTHESIS SPACE

3.3 Processing Algorithm

3.3.1 Overview

This section describes the computational algorithm used to generate likelihood values in the 4-dimensional target hypothesis space. In simplest terms, the core of this algorithm applied to a single radar DIR is a loop that, for each latitude and longitude cell within the geographic area of interest, and for each v_{lat} and v_{lon} velocity bin within that cell, determines all relevant range-Doppler-azimuth cells within the DIR, weights their value by the spreading function, and sums them together. However, this algorithm is not the most efficient way to calculate the quantity of interest, nor does it always produce the best results. Here are some of our departures from the basic process:

1. We have added an antialiasing step. We typically produce a grid of velocity bins at each location of size $N_v \times N_v$ with a minimum and maximum velocity V_{min} , and V_{max} in each direction. For the antialiasing step, the likelihoods are first computed for a much larger array of velocity bins $N_{vaa} \gg N_v$ with the same V_{min} and V_{max} . The effect of this step is to oversample the velocities. This oversampled array is then smoothed and downsampled before the final values are written out.
2. We really do not need to compute a separate likelihood value for each velocity bin. Recall that the velocities from a particular range-Doppler cell always form a linear stripe in the v_{lat}, v_{lon} velocity plane. Therefore, our algorithm converts overlapping groups of range-Doppler bins into a likelihood in ground velocity coordinates and we then project each of these values along a line normal to the radar steer direction.
3. We found that, when converting a range of sensor space measurements to ground space coordinates, a weighted sum did not always produce the best results. This is because a weighted sum is not a robust estimate of the target likelihood. In our case, the weighted sum tended to obscure strong target returns that happened to fall on the edge between target-hypothesis resolution cells. Drawing on our experience with OTH radar tracking and detection, we used an order statistic to compute the target-hypothesis likelihood value instead of a weighted sum. This way, a strong target return maintains its prominence in the transformation from sensor-to-target space. We found the maximum to be an effective order statistic for this experiment. In some cases, strong clutter returns may dominate a maximum-based algorithm, but we felt safe in using the maximum because the effects of clutter had already been removed in the normalization process prior to fusion.

3.3.2 Algorithm

The following pseudo-code describes the sensor-to-target-space mapping algorithm. We start with a sequence of radar dwells from both radars ordered in time.

Input: A sequence of normalized, radar dwell measurements in log-likelihood form, $llrs(r,d,a,s,t)$, where
r is the range index, ranges from r_{min} to r_{max}
d is the Doppler index, ranges from d_{min} to d_{max}
a is the azimuth index, ranges from a_{min} to a_{max}
s is the sensor (for now, either Texas or Virginia)

t is the time the dwell was taken, ranges from t_{\min} to t_{\max}
Output: A sequence of target-hypothesis space data sets, $llrt(lat, lon, vlat, vlon, t)$, that span the rectangular region bounded by lat_{\min} , lat_{\max} , lon_{\min} , and lon_{\max} . Each geographic bin is d_{lat} by d_{lon} in size.

lat is the center latitude for this hypothesis,

lon is the center longitude for this hypothesis,

v_{lat} is the latitudinal velocity component for this hypothesis

v_{lon} is the longitudinal velocity component for this hypothesis

t is the time of this hypothesis.

Additional definitions:

min_v : the minimum possible target velocity

max_v : the maximum possible target velocity

MINLLR: the minimum value a log-likelihood ratio could ever assume

OSRATIO: that ratio of oversampling

$tmpv[N_v * OSRATIO]$: temporary storage for velocities

Anti-aliased mapping algorithm

// Step through time

for $t = t_{\min}$ to t_{\max} , do

// Loop over latitude

for $lat = lat_{\min}$ to lat_{\max} , do

// Loop over longitude

for $lon = lon_{\min}$ to lon_{\max} , do

cell_list = list of range/azimuth cells that contribute
to this lat/lon cell;

// Loop over all Doppler cells, converting to ground velocity

// and filling array tmpv

vindex = 0;

for $d = d_{\min}$ to d_{\max} , do

v = Doppler d converted to ground velocity;

max = MINLLR;

for each range/azimuth pair (r,a) in cell_list, do

if $llrs(r,d,a,s,t) > max$

max = $llrs(r,d,a,s,t)$;

end

// Fill in oversampled values of velocity

for $i = 0$ to OSRATIO do

tmpv[vindex + i] = max;

end

vindex++;

end

smooth array tmpv with a Gaussian weighted average;

// Loop over all velocity cells

for $v_{lat} = min_v$ to max_v , do

for $v_{lon} = min_v$ to max_v , do

$v = (v_{lat}, v_{lon})$ projected onto a line in the
direction of the radar steer;

$llrt(lat, lon, v_{lat}, v_{lon}) = v$;

end

end

end

end

end

3.4 Statistical Modeling

3.4.1 Error Correlation Study

3.4.1.1 Approach

In order to validate the assumption that target information is uncorrelated between two distant OTH sensors, we have studied the correlation between position errors in the Virginia and Texas radar sites. There are actually two types of correlation that are relevant to this situation: amplitude and position. Amplitude correlation is widely studied in OTH radars. Due to the dependence of short-period radar echo variation on ionospheric phenomena, we are confident that radar echoes propagated through two widely disparate ionospheres are uncorrelated in amplitude. However, to our knowledge, there has been no equivalent analysis of the positional errors between two actual radar sites.

Studying all possible cross-correlations between sequences of position vectors can be complex. Since we are not sure along which axis the errors might be correlated, a complete investigation would need to generate cross-correlation sequences between all possible projection axes from each radar. Figure 3-9 illustrates this problem. The question we must answer is:

For a pair of sensors X and Y, when the position errors are projected along any two directions θ_X and θ_Y , are the errors uncorrelated for all time lags τ ?




FIGURE 3-9 THE PROBLEM OF DETERMINING CORRELATION BETWEEN TWO VECTOR POSITION SEQUENCES

Thus, we are examining a 3-dimensional array of correlation values in the space of the variables θ_X , θ_Y , and τ . Visualizing this space can be somewhat of a problem, so we have taken slices along various 2-dimensional planes to present the results. These planes will be determined by the data itself. For each set of vector position errors, we have determined the major and minor axes of the covariance ellipses as illustrated in Figure 3-9. We have reasoned that, if a correlation between errors exists, it will most likely reveal itself along one of these axes.

The mathematical tool we will use to analyze correlation is the cross-spectral density (CSD). Simply put, the CSD is the Discrete Fourier Transform of the cross-correlation between two

sequences. In these examples, we have used Welch's average periodogram method to compute the cross spectrum.

The analysis procedure used is illustrated in Figure 3-10.

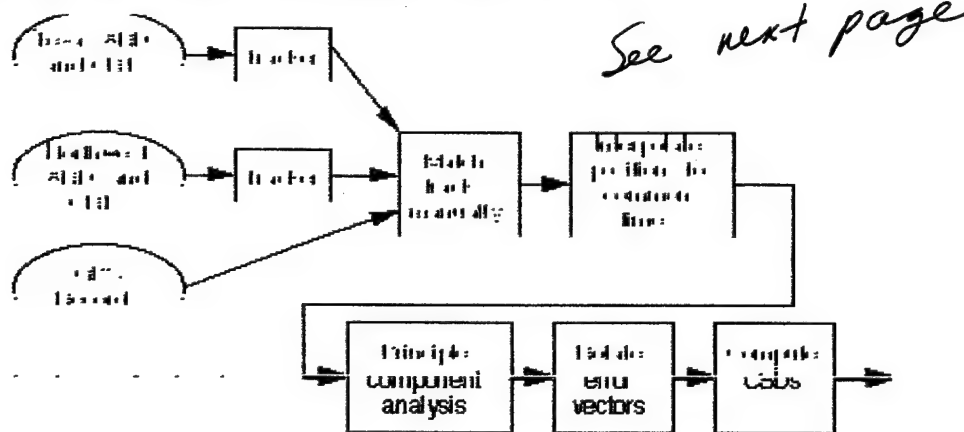


FIGURE 3-10 PROCEDURE USED TO ANALYZE THE CORRELATION BETWEEN OTH RADAR POSITION ERRORS

The steps are as follows:

1. Track Texas and Virginia data separately.
2. Compare tracks to GPS data and manually select the correct tracks that correspond to the control target.
3. Interpolate both sets of tracks onto a common time reference as determined by the GPS data.
4. Subtract the tracked position from the GPS position to obtain the position errors in latitude and longitude.
5. Perform principle component analyses on each set of errors separately.
6. Project each set of error vectors onto one of the principal directions for that data set.
7. Compute the CSD for the projected error vectors.

3.4.1.2 Data Sets

Our original intent for this study was to use the Caribbean radar system (CERAPS) tracks as ground truth. We received a set of these tracks from NRaD, acquired during a period of simultaneous ROTHRA VA/TX operation in an area of potential overlap. Unfortunately, after some analysis, we determined that the CERAPS tracks were not located in an area of overlapping coverage. As an alternative, we selected the 31 May 95 P3 Flights. This data set consisted of a P3 aircraft flying a 360-degree course while the TX and VA radar dwell illumination region (DIRs) were moved to provide overlapping coverage. We would have preferred the diversity of targets that the CERAPS data would allow, but the P3 set provides several hours of quality data

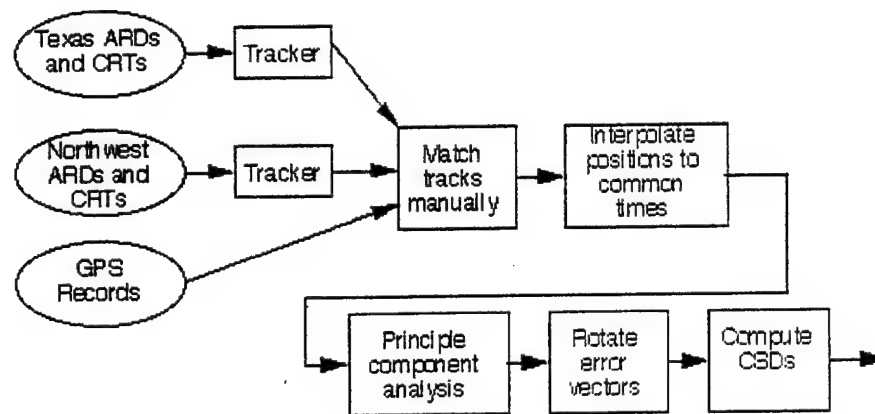


FIGURE 3-10.

on a precisely geolocated target. Figure 3-11 shows the course of the P3 from GPS, and the tracks from the individual radars as extracted from the range-Doppler data.




FIGURE 3-11 A SINGLE P3 TARGET TRACK AS REPORTED BY GPS AND AS TRACKED BY TEXAS AND VIRGINIA RADARS

Note that the segments of tracks shown in blue and yellow in this plot match the shape of the red GPS track closely but are shifted in bulk. This is most likely due to incorrect mode assumptions in putting the tracks to ground.

3.4.1.3 Results

Figure 3-12 shows the distribution of errors for these data sets plotted in latitude and longitude coordinates.

While this presentation shows the general clustering of errors, there is not consistency between the two radar errors. It is difficult to see a common trend between the two. Figure 3-13 also depicts the errors plotted in latitude and longitude, but the centroid of each error cluster has been removed. This figure shows the similarity in variance between the two error clusters, but there is still no trend evident. The nature of the position error is not evident until we plot the errors in terms of radial and azimuthal variation. Figure 3-14 shows the same error clusters as the previous figures, except each error point is plotted in coordinates of range and bearing referenced to the appropriate radar site. Now we can see that the errors are clustered around discrete, predictable locations in range with some random variation in the azimuthal direction. The discrete range error bins most likely correspond to particular propagation modes that were incorrectly identified in the slant-to-ground coordinate transformation.




FIGURE 3-12 POSITION ERRORS FROM TRACKED P3 TARGET PLOTTED IN LATITUDE- LONGITUDE SPACE




FIGURE 3-13 POSITION ERRORS FROM THE ROTH R RADAR WITH THE CENTROID OF EACH ERROR CLUSTER REMOVED.




FIGURE 3-14 CENTRALIZED ERRORS PLOTTED IN RANGE AND AZIMUTH COORDINATES.

Given the errors, the next step in the process is to extract the principal components for each cluster, which we have done by computing the eigenvectors of the covariance matrix for each cluster. This determines the major and minor axes of the error ellipse for each group. Figure 3-15 shows the same error clusters plotted in these coordinates.




FIGURE 3-15 POSITION ERROR CLUSTERS PLOTTED ALONG PRINCIPAL COMPONENTS OF THE INDIVIDUAL ERROR CLUSTERS.

The next step in the error analysis is to project the errors onto these axes and compute the cross-spectral density between the two radars. A strong peak in the cross-spectral density will

indicate correlation between the two error sequences at the frequency where the peak appears. Figure 3-16 shows the four cross-spectral densities that result from comparing the major and minor axes of the two radars.




FIGURE 3-16 A COMPARISON OF CROSS-SPECTRAL DENSITIES BETWEEN MAJOR AND MINOR AXES OF TEXAS AND VIRGINIA POSITION ERROR SEQUENCES.

In this figure, the CSDs have been normalized so that the peak value over all data sets is 1.0. Note that there is very little significant difference between the plots. All these CSDs exhibit a peak at about 0.0008 Hz and then drop off sharply. All other correlation values are so low that we can safely say there is no other correlation between the position error sequences. Of course, there is no guarantee that these are the only directions along which a correlation might exist. To determine if there are some directions other than the principal axes that might be correlated, we have computed the CSD while rotating one set of errors and holding the other fixed. Figure 3-17 shows the effect of rotating the Virginia error set and computing the CSD against the minor axis of the Texas radar. Although we studied various combinations, we chose to display this data set because it had the most variation. Again, we can see that for most rotations, there is a peak at approximately 0.0008 Hz and the rest of the data set is effectively uncorrelated. The strongest correlation peak occurs at approximately 180 degrees, which happens to be the approximate steer of the radar during this period.

Figure 3-18 shows the same experiment for the Texas radar. In this plot, the Texas error set has been rotated through 180 degrees and the CSD computed against the minor axis of the Virginia radar. Again, we see the peak only at 0.0008 Hz, with the strongest point along the steer of the radar.




FIGURE 3-17 CSDs OF VARIOUS VIRGINIA ERROR COMPONENTS COMPARED TO THE MINOR AXIS OF THE TEXAS ERRORS




FIGURE 3-18 CSDs OF VARIOUS TEXAS ERROR COMPONENTS COMPARED TO THE MINOR AXIS OF THE VIRGINIA ERRORS

3.4.1.4 Analysis

It is evident that there is no correlation between the two data sets at the frequencies that are relevant to our work. For example, a typical DPA tracking segment might be 5 or 10 minutes long at the most. This correlation lag corresponds to frequencies of 0.0017 Hz and greater. The CSD in this frequency range is below -50 dB for all cases. We can consider this amount of correlation unmeasurable. Any correlation in this range is purely an artifact of the computation. However, from an academic perspective, we would like to account for the correlation that exists at lower frequencies, *i.e.*, the peak at about 0.0008 Hz. This frequency corresponds to a correlation lag of about 20 minutes, which happens to be the length of the straight segments that the control aircraft was flying between turns. Thus, it appears that the position errors are correlated to each other, and to the target motion. We can also observe from the correlation data that the strongest correlations at this frequency exist along the steer direction of radars. This suggests that the peak may be related to a phenomenon that affects range error to a much greater degree than azimuthal error.

One explanation that accounts for these observations is that the correlation peak results from the tracking process itself. Note in Figure 3-11 that the tracks tend to jump modes at the turns. This is because the tracker frequently drops the target during the turn. During straight segments, the mode on which the track is being developed may be fading, but the tracker extends the track anyway because the early track history lends credibility to later detections, even if they are not coming from the currently dominant mode. However, when the target turns sharply, the tracker can no longer extend the existing track. When it begins a new track, it is likely to be on the mode that is dominant when the track is initiated. Thus, we have groups of errors, each of which correspond to a particular mode of propagation, and last exactly as long as the target continues in a straight line. The impact of this fact on our current research is minimal, but the message to the overall OTH tracking community should be to beware of unexpected correlations in the data that result from a combination of target behavior and ionospheric effects. This is particularly relevant to tracking stages that form long tracks by linking the short segments produced by detection algorithms such as DPA.

3.4.2 Spatial Uncertainty Model

The sensor-to-target-space mapping requires knowledge of the probability density function $p(z_k^{target} | z_{ki}^{sensor}, \theta_k)$. This function describes the probability that a given cell in the sensor measurement space (a range-Doppler cell) could contribute to a particular cell in ground-coordinates. Determining this function precisely could be a major task in itself, requiring considerable ground truth information in the form of air traffic control tracks, on-board GPS sensors, or radar repeater beacons. We have made extensive use of a study done for the Virginia system based on a number of GPS-logged flights of Piper Aztec and Queen Air aircraft in the Caribbean. A histogram of the error of several hundred position reports was constructed from this data by Navy personnel (Figure 3-19). This histogram shows two sources of position error. First, there is a predominant central peak that surrounds zero in this distribution. Second, there are several smaller peaks at outlying positions. The central peak corresponds to random distribution of position errors when the correct mode assumption is made for slant-to-ground transformation. The outlying peaks represent incorrect mode assumptions in the coordinate transformation.

For the purposes of this proof-of-concept project, we have only attempted to model the central mode of this distribution. It is perfectly reasonable to model the additional peaks as well, and to include the mode assumption uncertainty in the target-space mapping. If we take the mode uncertainty into account, then the fusion process and the ionospheric modeling process can work in tandem. The ionospheric model can produce a "best guess" model with uncertainty parameters. The fusion process can then combine not just multiple radar returns, but also multiple mode returns, into the same target hypothesis space.




FIGURE 3-19 POSITION ERROR HISTOGRAM FROM PIPER/QUEEN AIR EXPERIMENTS

3.5 Performance Prediction

If the coordinate registration procedures always gave perfect results, there is no question that our fusion approach would improve the detection sensitivity of a multi-site (or even multi-mode) OTH radar system. Unfortunately, this is seldom the case. OTH radar geolocation performance is notoriously unreliable. Figure 3-19 illustrates this point quite well. Even when we disregard the outlying locations arising from incorrect mode identification, there is a variance of several nautical miles in the central lobe of the distribution. We must ask ourselves if the fusion process we are developing can actually improve the detectability of small targets in the face of realistic coordinate registration errors. If not, there is no reason to proceed with further experiments. To provide an answer to this question we have undertaken the following performance prediction study.

Given the multi-radar detection procedure and the statistical models defined in the previous section, we would like to predict the realistically achievable improvement in radar performance. In particular, we would like to determine the effect of uncertainty in coordinate transformation (position errors). To address this issue, we have devised a Monte Carlo simulation based on

actual data collected from the Texas and Virginia radars. We have chosen to approach the problem this way due to the lack of data for actual control targets. Although there is considerable data with corresponding air traffic control tracks in one radar or the other, there is little or no overlapping coverage in these areas for a statistically significant variety of ionospheric conditions. Our Monte Carlo simulation approach allows us to control the amount of uncertainty in the registration procedure so that we can observe the effects on the fused SNR over a track.

Figure 3-20 shows the procedure we have used to generate the performance prediction results.

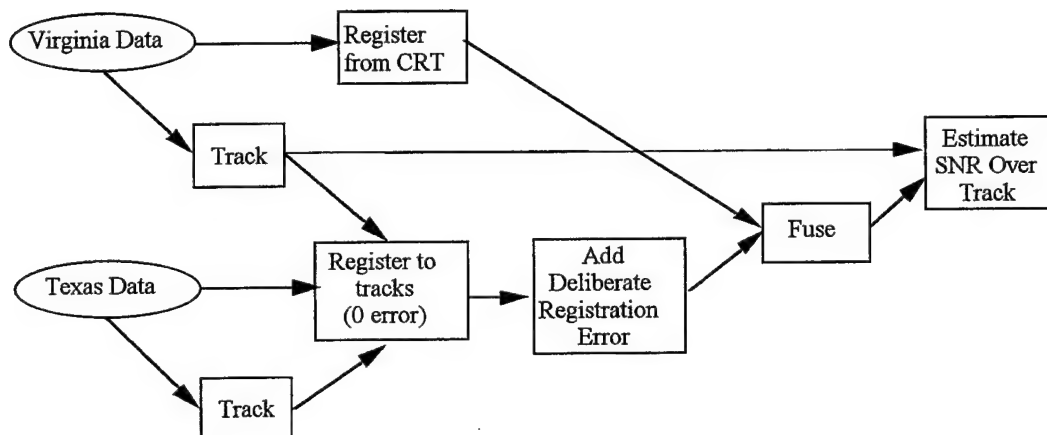


FIGURE 3-20 PROCEDURE USED TO PREDICT FUSION PERFORMANCE IN THE PRESENCE OF COORDINATE REGISTRATION ERROR

We start with raw data from each radar taken from an area of overlapping coverage. We then apply the SRI tracker to the individual data sets. We must run the tracker to produce targets that can be matched between the two data sets. These matched targets are then used as the reference for artificially perfect coordinate registration. Like our fusion approach, the SRI tracker also uses the DPA tracking procedure, so the tracks also produce an estimate of how the fusion process works in the single-radar case. We then send the Virginia radar tracks to ground using the ROTHr coordinate registration table (the choice of Virginia as the reference was arbitrary). The Texas tracks are then matched by hand to the Virginia tracks to determine a Texas coordinate registration that produces zero registration error compared to the Virginia data. To this coordinate registration, we then add a random offset in the range direction. The random offsets are samples taken from a Gaussian distribution with a known variance. We then apply the sensor-to-target space mapping procedure to both data sets, using the ROTHr CRT in the Texas case and the randomly offset CRT in the Virginia case. Finally, we accumulate likelihoods in the known target positions to produce a cumulative target SNR over the length of the known track. We perform this procedure several times for each value of position variance until we obtain a stable average of the track SNR.

The choice of Gaussian distributed position offsets is based on the distribution shown in Figure 3-19. We have chosen to model only the central lobe of this distribution and not to reproduce any of the second-order effects, such as the skew due to "ROTHr long" errors or the small irregularities in the shape of the curve. Future studies should investigate the result of position bias in addition to random errors. A different random position offset is added to each radar dwell, and the offsets are statistically independent from dwell to dwell. Our justification

for this model comes from the results of the previous section where we have shown that position errors are uncorrelated over inter-dwell periods.

Figure 3-21 shows the results of this study. The pink line represents the cumulative target SNR from the Texas radar alone, and the yellow line shows the SNR from the Virginia radar. The blue line shows the results of the Monte Carlo simulation at various standard deviations (σ) of position error. At $\sigma = 0$, the SNRs from the individual radars add perfectly. Of course, this is an unrealistic situation, but it serves as a check on our results. From this figure we can see that the fused SNR is greater than either radar taken alone, even for σ values of 40 nautical miles (nmi). From Figure 3-19, we judge the typical ROTH range registration σ to be just over 20 nmi. At this value of σ on the plot of Figure 3-21, we can see a gain of approximately 5.5 dB. Thus, we believe that significant reduction in minimum detectable SNR is possible with realistic ROTH coordinate registration accuracy.

Note that we have only modeled errors that occur when the correct mode is identified by the CR procedure. If modes are not identified correctly, then a model for this behavior should also be integrated into the mapping procedure and not left as outlying errors. Also, we have not attempted to model the behavior of the tracking process itself. We have assumed that the tracker will behave at least as well as it did in the single-sensor classes. In reality, since we are using a joint tracking and detection algorithm, the tracking performance will vary with the quality of the data. However, this study lets us isolate the effect of coordinate registration accuracy on the SNR of the target, tracking issues aside.




FIGURE 3-21 PLOT OF FUSION SNR GAIN VS. POSITION ERROR STANDARD DEVIATION

4 RESULTS

We have applied the fusion/tracking algorithm to a variety of data sets. The challenge lay in finding interesting control target behavior during a period of continuous, overlapping coverage from both radars. The first example in this section tracks the P3 control target through the trajectory shown in Figure 4-1. This example demonstrates the viability of the technique, and shows that targets can be tracked continuously through mild maneuvers. However it does not demonstrate the unique ability of this approach to track targets through maneuvers that take them

through zero Doppler blind speed in one radar or the other. In order to demonstrate this capability, we injected a realistic target undergoing a 180-degree turn into an actual radar measurement set. The trajectory of this simulated target is shown in Figure 4-2.

We have applied the fusion approach to several sets of overlapping data from the Texas and Virginia radar sites. We present results here drawn from the 31 May 1995 P3 flight, and also have injected a target into real data to simulate a target in a tight turn.

Figure 4-3 illustrates results of applying the fusion/DPA tracking algorithm to the combined Virginia and Texas data set. In this figure, the results are presented on a grid spanning the geographic location between 17.0 N - 18.5 N latitude and 73.75W - 75.5W longitude, with each cell covering 0.25 deg in latitude and longitude. The frames in this image highlight the maximum velocity hypothesis score for the geographic region. Thus, the peaks in these images are maximum over the entire tracking region. Within the cell containing the maximum score, the vector from the center of the cell to the largest peak represents the target velocity and the direction of motion.

These snapshots were taken at intervals of 2, 13, 21, 37, 49, 51, 52, 53, 56, 58, 79, 92 and 104 dwells into the tracking process, where the number of dwells is measured cumulatively over the two radars starting at 22:40 on 31 May 1995. Note that the peaks in these snapshots are quite elongated in a direction normal to the steer of the Texas radar. This is due to the relative revisit rates of the two radars. In this data set (31 May 1995), the Texas radar was being operated in a staring mode with revisit rates as short as every 15 seconds. However, the Virginia radar was operated in a standard manner, meaning the revisit rate was close to 30 seconds. Since the steers of the two radars are more or less orthogonal, the resulting estimates are made with nearly twice as much information along the range axis of the Texas radar as there is along the range axis of the Virginia radar. Because we are assuming that the dwells are statistically uncorrelated (an assumption borne out by analyses reported earlier), the variance of the velocity estimate along the Texas range axis is much less than the variance along the Virginia axis. Thus, the peaks shown in Figure 4-3 are elongated. The target is seen flying in the northeasterly direction at approximately 24.5 deg to the horizontal. At 22:30 the target makes a course change and continues to fly in the 53.6 deg to the horizontal. This course change is observed as a change in the location of the maximum score between dwells 52 and 53 in Figure 4-3. In addition, it is also apparent that the target travels farther north than east after dwell 53.




FIGURE 4-1 TRAJECTORY OF THE P3 AIRCRAFT DURING THE PERIOD OF THE TRACK

At this point, it is apparent that the real advantage of the fusion process and the DPA tracking algorithm is in tracking a target through a turn. However, there is insufficient data available with both the Texas and Virginia radars looking at the same region where a target is making a sharper turn. We have overcome this lack of data problem by simulating a target flying in the overlap area of the Texas and Virginia radars and injecting this simulated target in the corresponding data. We have the simulated target heading northeast at 45 deg to the horizontal at 22:45. After 6 minutes, the target executes a course change of 180 deg. and heads in the southwesterly direction for the remainder of the tracking cycle. Two turn accelerations are considered: in Figure 4-4, the target makes the half-turn in 3 minutes, and in Figure 4-5, the same turn is accomplished in 2 minutes. In both cases, this approach is able to track the target through the sharp turn. The target passes through the clutter for both the radars, but at different times, and hence the tracking process is able to reliably follow the target through the turn. With additional radars looking at the same region, the tracking capability will be further enhanced.




FIGURE 4-2 TRAJECTORY OF SIMULATED TARGET MAKING 180-DEGREE TURN

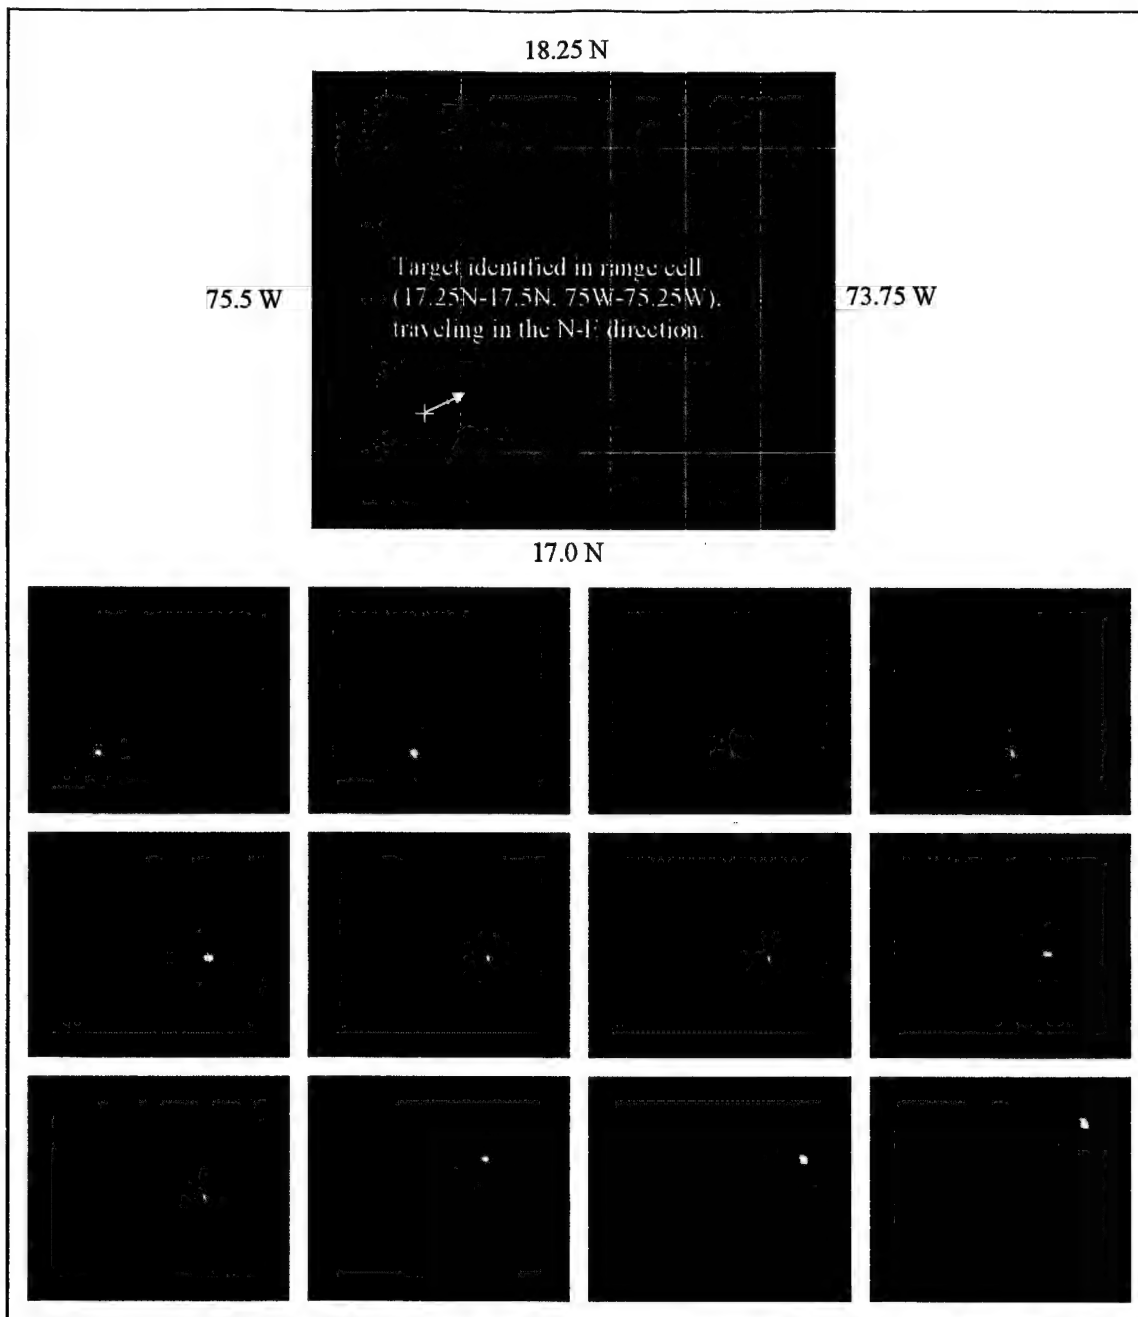


FIGURE 4-3 RESULTS OF TRACKING A REAL TARGET OBTAINED BY FUSING INFORMATION FROM VIRGINIA AND TEXAS RADARS BETWEEN 22:20 AND 22:40 ON 31 MAY 1995.




FIGURE 4-4 RESULTS OF TRACKING AN INJECTED TARGET OBTAINED BY FUSING INFORMATION FROM VIRGINIA AND TEXAS RADARS BETWEEN 22:45 AND 23:15 ON 31 MAY 1995. THE TARGET EXECUTES A 180 DEGREE TURN IN 3 MINUTES.




FIGURE 4-5 RESULTS OF TRACKING AN INJECTED TARGET OBTAINED BY FUSING INFORMATION FROM VIRGINIA AND TEXAS RADARS BETWEEN 22:45 AND 23:15 ON 31 MAY 1995. THE TARGET EXECUTES A 180 DEGREE TURN IN 2 MINUTES.

5 CONCLUSIONS AND RECOMMENDATIONS

5.1 Overview

Detection of aircraft targets in operational radar systems like ROTHr (and Jindalee) is traditionally performed in radar coordinates, and candidate tracks are started on each of the images of a target received by the radar receiver; that is, on each ionospheric raypath mode perceived by the radar. Separate radars surveying the same area can observe the same targets, and each develops its own set of candidate tracks. While a single mode (usually the low-low-F2 mode) is likely to possess higher average energy than other modes, the output of the signal processor for any given CIT can exhibit the highest signal-to-noise ratio for any mode present, owing to the path-interference and polarization-rotation fading suffered by all ionospheric paths. The tracker must take this into consideration in setting association gates and in assessing cumulative track scores (CLLR). Candidate tracks thus started are then transformed to earth coordinates before fusion with tracks derived from other radars is attempted.

In contrast, this project, which developed an integrated Multiple Radar Detection and Track (MRDT) scheme, endeavors to show how data from targets observed by two radars can be fused by starting tracks on data which has already been converted to earth coordinates, and in which fusion of the data from the two radars is started before tracking is established. The advantages have already been discussed, but the resultant 4-D target data, consisting of amplitude vs. vector-velocity in each geographic "pixel," would serve to alert the tracker to maneuvers of the target before any changes in geographic coordinates of the target are sensed. This feature, alone, should improve the ROTHr tracking accuracy and sensitivity through significant shrinkage of maneuver association gates in the tracker. This radical departure from the usual ROTHr method of track establishment may seem to offer difficult obstacles to effecting the process in ROTHr, but SRI believes that it is plausible to perform the necessary trials of the technique in ROTHr. This section deals with issues related to this transition of the MRDT technology to the ROTHr system and demonstration of the technology in the ROTHr Enhancement Demonstration System (REDS). Transition of MRDT technology to ROTHr involves a number of complex issues relating to architecture and connectivity, both present and future, as well as to future decisions regarding which research and development enhancements are to be transitioned. These issues will be discussed after a plan for demonstration of the technology in REDS is presented.

5.2 Performance Expectations

The illumination of a common area by both ROTHr-Virginia and ROTHr-Texas is a commonplace characteristic of OTH radar surveillance in the CD mission. Under all but rare situations, the illumination of this common area is on differing frequencies, for the two radars, and using different waveforms, *e.g.*, WRF, bandwidth, and CIT. In addition, the scan policies for the two radars are independent, and so illumination by the two transmitters is not likely to be simultaneous, meaning that fusion of the detection or track data from the two radars requires temporal interpolation as well as frequency/waveform normalization. Thus, we may consider that the two radars' echo signals from targets of interest in the common-illumination region are statistically independent of each other. More importantly, the major sources of "noise" in detection of target echo signals, clutter returns, are also mutually independent in the two-radar case.

By merging the signal-processed data from the two radars before tracking has been established, as is done in MRDT, the overall fading intervals should be halved, statistically, which translates to an improved detection sensitivity against noise and clutter. In the worst case, when a target is moving such that its echo falls into the ground clutter for one radar, it cannot simultaneously do so for the other, and the improvement in mean SNR, relative to that of the better of the two radars, is zero. However, there is another improvement possible, for the case when the target moves with varying Doppler frequencies (not solely attributable to radar operating frequency changes).

A track-while-scan radar like ROTHr revisits the targets in a DIR every 30-50 seconds under normal circumstances, and variations in target Doppler caused by maneuvering are perceived by the radar as step-wise Doppler jumps, or discontinuities. Kalman filters in the tracking function might exploit these Doppler jumps in anticipating range or azimuth (or latitude/longitude) association gate width changes, but the ambiguities available reduce the value of such a process. The MRDT tracking function, which tracks target echoes in four-dimensional space (latitude, longitude, and 2-D velocity [speed and course]), is sensitive to changes in any of these dimensions, and can accommodate association gates set to anticipate variance in any or all of these dimensions. By tailoring the shape of the association volume to anticipate large course variance with relatively small speed, latitude, or longitude variances, the tracking process can accommodate realistic aircraft maneuvers with a net tracking sensitivity that is undiminished relative to standard 3-D trackers working against nonmaneuvering targets.

The combination of these benefits has been coarsely estimated at 6 dB equivalent SNR improvement, through the experimentation described earlier in this report. Separate from these benefits, which are quantifiable in terms of equivalent SNR improvement or equivalent target RCS sensitivity improvement, there is the saving in detection/tracking operator actions required in the surveillance overlap region, since a single tracker is used to develop tracks and updates in this region. If the overlap area is very small, owing to the nature of the ROTHr tasking at the moment, then the benefit is small, but in a large-overlap situation, the manpower savings can be significant.

In a radar system like JORN (or OTH-B) in which the overlap regions, although not large, represent seams between two radars' coverages, MRDT processing can reduce the inter-radar handoff effort required, which otherwise is more operator-intensive than tracking in the clear areas served by a single radar. The manpower savings realized in track handoff across radar coverage seams, however, is expected to be modest, compared with the effort saved in repairing tracks of maneuvering flights, especially in CD missions.

5.3 REDS Implementation

MRDT technology can be demonstrated in REDS in the playback mode with the existing configuration, or in real time with the advent of the communications upgrade that will make complex ARD maps available to the DEC Alpha 8200 via DECNET (see Table 1 in Section 5.4.1 for Alpha 8200 description). LOGAMP data from both the Virginia and Texas systems will be accessible by the Alpha 8200 in REDS on playback tape. REDS implementation of MRDT would be done for a limited number of DIRs consistent with the processing requirements

of the two-step process of sensor ARD to geographic (latitude-longitude) pixel mapping and the Dynamic Programming Algorithm (DPA) tracking discussed below.

5.3.1 Playback Demonstration

Initial demonstrations of MRDT technology could be conducted in REDS using a merged playback tape created from data recorded from ROTH-R-Texas and ROTH-R-Virginia. This tape would be formatted on 4 mm DAT from two ROTH-R TK 90 tapes using the VAX 6430 and DECNET connection to the Alpha 8200. This tape would provide LOGAMP data as well as coordinate registration tables (CRTs) from both radars.

5.3.2 Real-Time Demonstration

Real-time demonstrations of MRDT could be conducted in REDS after the communications upgrade is in place. This upgrade will allow delivery of complex ARD maps from both Virginia and Texas to the Alpha 8200 via DECNET from the OCC (or possibly from the communications processor, directly). CRTs from both radars can be delivered from the OCC via DECNET; however, the conversion of complex ARD maps to LOGAMP data, as well as any normalization necessary, would have to be done in the Alpha or other REDS assets.

5.3.3 Processing Requirements

Estimates of the processing requirements that will be imposed by MRDT in a typical implementation such as REDS have been made for memory and speed. The first issue to be addressed is random access storage. If the process is to run without interruption, the entire procedure should take place in core memory, without the necessity of swapping in virtual memory pages. Thus, the memory requirements should not exceed the available RAM on the processing computer. Recall that the fusion processing consists of two steps: mapping and tracking. If these two steps are combined into a single process, then it is only necessary to keep the cumulative log-likelihood ratio output of the DPA tracking step in memory at any one time. The size memory required to accommodate this storage depends on the spatial and velocity resolutions to be achieved. First, assume that each DIR contributes to approximately 1024 locations on the ground (32 latitude cells by 32 longitude cells). For each DIR, an approximate 256 x 256 grid of velocity hypotheses must be computed. If this is done for each of three mode hypotheses, then each DIR requires storage of approximately 800 MB. Even for 16 DIRs, this is within the capabilities of modern, server-class workstations. In the case of the REDS Alpha 8200 the coverage would be limited to two DIRs.

The second issue is processing time. The fusion process should be able to compute cumulative scores for all ground locations in the time it takes for one sweep through the DIR set. In other words, the time required to process one DIR should be no greater than the revisit time, usually 30 to 40 seconds. The sensor-to-target-space mapping can be viewed as a sort of convolution. For each geographic/velocity hypothesis in ground coordinates, a set of range-Doppler cells from the sensor measurement space must be multiplied and summed together. The amplitudes (actually log-likelihood ratios) from each range-Doppler cell within the set must be multiplied by a scale factor (log of a conditional probability density function) and then accumulated in a target hypothesis bin. The size of the relevant set of range-Doppler cells depends on the

uncertainty in registration, but experiments suggest that 16 range-Doppler cells per target-space hypothesis is a reasonable estimate. For the 32x32x256x256 hypothesis space described above, this amounts to about 1 billion floating point operations every 30 seconds, or about 33 million floating point operations per second (Mflops) per DIR. This is the computational requirement of the mapping stage. The detailed computational requirements for the tracking stage are well documented in other reports [Shaw et al., 1996] dealing with the SRI/Raytheon integrated tracker (ITRACK), but for a transition search area of 3 cells in each direction, the processing is estimated to be about 6 billion floating point operations every 30 seconds, or 200 Mflops. This leads to a combined total of 233 Mflops per DIR for the entire processing stream. This processing need not be done on a single processor since the parallel nature of the DPA process would allow each DIR to be processed on a separate workstation. Thus, due to the highly parallelizable and pipelined nature of this processing, MRDT is well within the capabilities of modern workstations.

5.4 Architecture

The architecture necessary to support the MRDT computing requirements will soon be available in the REDS Alpha 8200 workstation described below.

5.4.1 DEC Alpha 8200

The Alpha Server 8200 is one of the highest-performing office systems in the industry, sporting a robust 64-bit architecture. Primarily used to support enterprise databases, complex simulations, and decision support applications, the Alpha Server 8200 offers an array of capabilities. Table 1 lists its specifications.

The Alpha Server 8200 is a planned replacement for the VAX 6430. It will be used to install Intergraph capability, replacing the current R2000 display consoles. The Alpha server will be available for software R&D and will maintain raytracing software (rehosted from ALPHA 2100) used in RTRT. It is connected to the Storage Works Cabinet, which provides for massive amounts of data storage.

5.4.2 Storage Works Cabinet

The Storage Works cabinet provides the Alpha server 8200 with over 50 gigabytes of storage capacity. It is designed to meet growing demands for capacity (storage can range into the terabytes) and performance (innovative cache and RAID capability that adapts to data throughput). Table 2 lists its specifications.

Table 1: Alpha Server 8200 Configuration

Alpha Server 8200 5/440 quad processor (expandable)
4 DEC chip 21164/437MHz processor module
• 2 gigabytes RAM

• (2) 4 megabyte backup cache
• Single session ASCII/ANSI PCTerm text terminal
• Open VMS-style keyboard
• (1) 600 megabyte CD-ROM drive
• (1) 2.1 gigabyte 3.5" SCSI drive
• Fast EtherWORKS PCI 10/100 network interface card
• 12-slot PCI plug-in unit for Alpha Server 8200 system cabinet
• EISA bridge option
• EISA-based high-speed synchronous communications controller
• (1) 4/8 gigabyte 4mm SCSI DAT tape drive
• Open VMS for Alpha OS distributed interactive user
• (supports 32 concurrent users)
• Digital FORTRAN for Open VMS Alpha
• MACRO-64 for Open VMS AXP
• DECset for Open VMS AXP

Table 2: DEC SW500 Cabinet Configuration

SW500 Departmental Cabinet
CPU Alpha 21164 processor, 266MHz
Storage works command console
Redundant single phase power option
Storage works 32 megabyte CI controller
(2) 10m CI-Bus cables
(6) 1m SCSI-2 cables
(12) 4.3 gigabytes SCSI fixed disks
(8) 4/8 gigabyte 4mm SCSI DAT tape drive

5.5 Transition to ROTHr

Prior to a decision to transition MRDT technology to ROTHr, a number of issues need to be resolved, among those being: multimode accommodation and mode linking, next-generation tracker compatibility, real-time raytracing compatibility, and re-hosted post-processor compatibility. Each of these issues is discussed below; however, at this time it is not possible to chart an unequivocal course for MRDT transition due to the uncertainties surrounding these issues.

5.5.1 Multimode Accommodation and Mode Linking

The preceding section on processing requirements for MRDT discussed RAM storage requirements for three propagation modes per DIR, even though the MRDT experiments to date have concentrated on exploiting only the dominant mode as determined from the ROTHr CRTs. The processing time given is for one mode only. There are two methods for doing the multimode processing. The first method would map all modes into the same hypothesis space on the ground. Each target-space hypothesis would accumulate LLRs from all possible mode hypotheses. In this method, the mapping-stage processing will scale linearly with the number of modes, but the tracking processing will remain the same. It will be important in this method to condition the contribution of each sensor measurement by the mode certainty. The probability of a target existing on the ground conditioned on a particular sensor measurement cannot exceed unity. If mapping is done for only one mode, the energy of a particular range-Doppler cell gets distributed only over the area on the ground to which it contributes. However, if mapping is done for multiple modes, energy must be distributed among the modes. If all modes are of equal certainty, this isn't an issue, but there is probably more confidence in the dominant mode than in the others. Therefore, that mode should get more weight in the final summation than less likely mode hypotheses. Weighting might be done in proportion to the predicted mode amplitude given in the CRT. Tracks generated from multiple modes could be linked in ground coordinates using a mode-linking approach similar to that currently in use by ROTHr.

The second alternative is to run through the whole process with a set of possible mode pairs from each radar. In this case, the entire processing requirements get multiplied by the product of the number of modes. In other words, if the three most prominent modes from each site were used, the processing would increase by a factor of nine. This process must be followed by a track merging stage in registered ground coordinates, which would replace the current mode linking approach used in ROTHr.

The choice of technique depends on the certainty in mode identification. If some of the modes are so uncertain that SNR is actually degraded by the smearing of energy over many cells, then it would be wise to use the second method. However, if there is equal spatial uncertainty for each mode hypothesis, it is probably better to use the first method because there is less processing required.

5.5.2 Next-Generation Tracker

A number of tracking enhancements are under consideration for transition to ROTHr, among those being the SRI/Raytheon ITRACK, the NRC/Raytheon MLANS, and the Raytheon Advanced Tracker. Only ITRACK uses a Dynamic Programming Algorithm (DPA) that is functionally equivalent to that used in MRDT. If either MLANS or the Raytheon Advanced Tracker is selected for transition implementation, then MRDT faces a significant integration obstacle. If ITRACK is selected, then it would appear that a significant portion of its DPA front end could be integrated with the processing required for MRDT.

5.5.3 Real-Time Raytracing

Two approaches are under consideration for utilizing real-time ray tracing (RTRT) to improve modeling and coordinate registration in ROTHr. One technique involves regenerating the CRT for each DIR (with all predicted modes) for each model update. The second involves

calculating slant-to-ground conversion factors for each slant track as they are required for mode linking. The first method is consistent with the current ROTHr mode-linking approach and with the MRDT mapping of ARD to geographic cells. The second approach, while possibly offering some advantages in propagation management assessment (PMA) architecture, may not be compatible with MRDT, although it might prove useful if the second method of MRDT processing (described above) is required. It would seem appropriate to try some combinations of RTRT and MRDT in REDS prior to committing to a transition plan.

5.5.4 Radar Post-Processor Re-Host

Current plans call for upgrading the Radar Post Processor (RPP) by rehosting the existing functionality onto a DEC Alpha platform in the OCC. The new RPP will receive complex ARD maps from the communications/site-separation processor and produce normalized ARD or LOGAMP maps. CFAR processing and peak picking are also to be done in the same module. Some, if not all, of the RPP functionality and processing will be either modified or replaced by MRDT processing, so it behooves the transition team to keep up with how the rehosted RPP development is proceeding.

5.6 Summary of Conclusions

SRI's conclusions based on this study are as follows:

- Multiple OTH radars with coverage in overlapping areas can be effectively fused at the raw data level.
- Combined tracking and fusion in the raw data, prior to the detection stage, enhances the detectability of small targets.
- The effectiveness of this technique depends on the accuracy of the coordinate registration system.
- Typical ROTHr coordinate registration accuracy produces a sensitivity gain of approximately 5 dB.
- The processing demands of this technique are within the capabilities of modern workstation computers.
- The effectiveness of this technique suggests that the coordinate registration process for all OTH radars could be recast as a statistical problem. There is a limit to the effectiveness of completely deterministic ionospheric modeling in the coordinate registration process. Studies have shown that the first- and second-order statistical properties of OTH coordinate registration errors are predictable, even if the exact error at any one time is not. Furthermore, we have shown that these errors are uncorrelated between radar sites. The current approach to coordinate registration, consisting of ionospheric modeling and inversion, is not a robust approach to the problem, i.e. it fails catastrophically when the model is incorrect. For example, mode identification is frequently incorrect. As an alternative, we suggest that the slant-ground coordinate transformation be viewed as a statistical process, so that random deviations from the ionospheric

model do not disproportionately affect the accuracy of the tracking process.

The decision on whether and how to transition the fusion process to ROTHr depends on a number of complex and interrelated issues. Some of these issues can be resolved by implementing MRDT in REDS and conducting experiments with playback or live data. Some of the issues require consultation with the ROTHr Program Office and with other organizations involved in the transition of certain ROTHr enhancements. The multimode issue must be resolved, and it is recommended that a REDS experiment be designed to evaluate the impact of considering multiple modes in MRDT implementation. The next-generation tracker issue requires discussions with those conducting the evaluation of competing tracker technologies.

REFERENCES

- Barniv, Y., "Dynamic Programming Solution for Detecting Dim Moving Targets," *IEEE Transactions on Aerospace and Electronic Systems*, AES-21, 1, January 1985.
- Barniv, Y., and O. Kella, "Dynamic Programming Solution for Detecting Dim Moving Targets. Part II: Analysis," *IEEE Transactions on Aerospace and Electronic Systems*, AES-23, 16, November 1987.
- Bar-Shalom, Y., "Extension of the Probabilistic Data Association Filter to Multi-Target Tracking," *Proc. Fifth Symp on Nonlinear Estimation*, pp. 16-21, San Diego, 1974
- Bar-Shalom, Y., and E. Tse, "Tracking in a Cluttered Environment with Probabilistic Data Association," *Automatica*, Vol 11, 1975.
- Blackman, S., *Multitarget Tracking with Radar Applications*, Artech House, Norwood, MA, 1986.
- Kewley, D. *Presentation to the Fifth U.S./Australian Joint Study Group*, September 1994
- Kurrien, T., and P. Milanfar, "OTH Radar Track Fusion," *Joint Data Fusion Symposium*, October 1994.
- Reid, D. B., "An Algorithm for Tracking Multiple Targets," *IEEE Trans. Automatic Control*, Vol. 24, No. 6, 1979.
- Sharf, L. L., *Statistical Signal Processing: Detection, Estimation and Time Series Analysis*, Addison Wesley Publishing Co., Reading, MA, 1991.
- Shaw, S., C. Woodworth, and J. Arnold, "OTH Radar Detection and Tracking Enhancements," SRI Final Report, Contract F30602-92-C-0110, SRI Project 4564, SRI International, Menlo Park, CA, April 1996.
- Wishner, R., F. Lanzinger, H. Mesiwalia, and S. Mori, "Advanced Techniques for Multi-Target Track Formation," Final Report to BMDATC, Advanced Information and Decision Systems, March 1981.
- Yssel, W., "OTH Sensor Fusion," NRaD White Paper, September 1994.

4.3 DATAFUSE AND ALICE IN AS SYSTEMS

Implementation issues for Datafuse at JFAS

Andrew J. Shellshear and Ian W. Dall

October 12, 1995

1 Introduction

Multiple target returns may result from looking at a coverage area with multiple radars, or from multimode returns on a single radar.

Here are some implementation issues of a multi-mode, multi-radar track fusion algorithm at JFAS for discussion by the JCT.

"Datafuse" is a program that fuses radar tracks caused by the same target, assigning ionospheric propagation modes to each OTHR track. It returns co-ordinate registered tracks (ie. in lat, long) that are a weighted combination of the radar tracks.

The aim of the implementation is to aid current fusion capabilities. Datafuse will automatically fuse tracks, and will assist operators in making fusion decisions.

The proposed implementation will be done in several steps. The eventual functionality will be described first, then the changes to current software necessary to achieve it, in several steps.

It is still not clear how much track information will be accessible from within the JFAS software. This report is written under the assumption (yes, Paul, I know) that external track sources such as Microwave radar tracks and ADS data will be accessible from inside the system. See Section 3.4 for details.

The displays shown in this document are intended to illustrate the functionality requirements. They have not been put into approved JFAS format.

2 Eventual Functionality

2.1 DTO and ECA Displays

The DTOs (Detection and Tracking Operators) and ECA (Environmental Conditions Advisor) will be responsible for supervising the fusion algorithm and doing manual fusions. By default, the DTOs will be able to do any multimode fusions, and the ECA all other types of fusion. The DTS (Detection and Tracking Supervisor) is responsible for selecting how much fusion is done and where it gets done. (See Section 2.4.)

2.1.1 ECA

The ECA is in charge of ensuring consistency in CR. They have access to the CR system and can make educated decisions that affect CR. Consequently, they get dumped with all the track fusion that gives CR feedback. This includes fusion of tracks with beacons, airplanes, microwave radar and ADS tracks. Once an OTH track has been fixed to a truth source, feedback to CR will happen automatically for as long as both sources are updated. The CR feedback may be turned on and off in the "who-to-tell" part of the "fuse tracks" window (see Section 2.12.2).

2.1.2 DTO

The DTOs will not be able to do any actions that effect C'R. by default. (See Section 2.5.) Hence all ability to set ionospheric heights from DTO displays will eventually be removed. DTOs will not be able to fuse OTH tracks with airplanes, beacons, Microwave radar or ADS tracks data. Each of these fusion process give unambiguous ground coordinates and provide feedback to the C'R.

The issue of fusing multi-moding tracks is somewhat more complex. It is possible to give feedback to C'R simply from multi-moding data (along the lines of "the ionospheric layers are too widely separated if these tracks are in these modes") but it is clearly impractical to disallow DTO operators from doing multi-moding, so there are two options. The first is to ignore any C'R feedback from DTO operators doing multi-moding but to provide feedback for the ECA doing multimoding. The second option is to allow DTO operators to provide C'R feedback from multimoding of this kind only.

The method of selecting modes for tracks will be greatly changed. The current A, B and C layers that the operator may assign to tracks will be discarded entirely, as they do not represent actual modes of propagation. Instead, operators will select from the available propagation modes. (For example "E E", "E F1", "F1 E", "F1 F1", etc.)

2.2 Fused Tracks

Any tracks that are fused by datafuse or operator action form a "track family". The tracks in the track family are not individually displayed. Instead a fused track representing the group is shown. (See Section 3.2.) In the displays that are in ground coordinates the fused track will be shown in lat. long or ground range, ground azimuth in a distinctive colour. In the displays that are in radar coordinates, the fused track position is converted to radar coordinates using the dominant mode, and displayed in a distinctive colour. (See Section 3.3)

All tracks in the track family are made invisible to the operator. (See Section 3.1)

The fused track is a weighted mean of the ground co-ordinates of each of the track family members.

The distinctive colour of fused tracks will be purple, replacing the current usage of purple tracks to represent tracks in the non-dominant mode.

2.3 Locking Tracks

It is possible that with multiple operators and the fusion program operating simultaneously, a track may be involved in several fusion operations at the same time (for example, if Datafuse decides to fuse several tracks just when an operator is doing a manual fusion on the same tracks). In order to avoid this a locking system will be active. When a track is entered into a complex validate, or is placed on an operator-assist queue, the tracks will be "locked" and any attempt to do a simple validate or manual fusion by the operator, or an operator-assisted or automatic fusion by Datafuse, will be disallowed. The track in question will not be included in the operation being done, and an error message will be displayed in the Message Window.

2.4 Automatic Fusion Window

A window will exist that allows the DTS (Detection and Tracking Supervisor) to select the degree of automation of the fusion process. (See Figure 1)

A bar shows levels of fusion confidence between 0 and 100%. Two sliders on the bar allow the operator to choose how confident datafuse must be before it can fuse tracks.

Automation of Fusion

Configure

DT Station Number

0%

1

40%

2

75%

3

100%

Fusion Confidence Level

Configure Fusion Automation

DT Station Number

Region

Region

Region

Region

Region

Region

Region

Region

Region

Region

Do fusion of:

☐ Multimode Tracks
☐ Beacons

☐ ADS Tracks
☐ Airlines

☐ Microwave Tracks

Associate after

2

 points on track

Change opinion after drop in confidence of

10

 %

Association memory time-constant

10

 minutes

Figure 1: Fusion Automation Window.

1. All track associations with confidence to the left of both sliders will be not considered for fusion.
2. All track associations with confidence between the two sliders will be operator-assisted. (see Section 2.11)
3. All track associations with confidence to the right of both sliders will be automatically fused. (see Section 2.2)

The "DT Station Number" tells which DT station this set of preferences is valid for. It may be changed from the Configure Button, which calls up the Configure Fusion Automation window.

2.5 Configure Fusion Automation

The configure button allows the DTS to screen the track-associations that operators see (either automatically or operator confirmed).

By clicking on the DT station number the DTS may change the DT station being configured. The Automatic Fusion Window will change to show the sliders for that DT station.

There is a full list of the current regions below the DT station number. The selected regions are those for which this DT station's preferences hold. All fusion done in these regions is supervised by this DT station.

Below is a set of fusion preferences. The DTS may assign what kind of fusion may be done by each operator. By default DTO's will only be able to do multi-mode, and the ECA will have full access to all kinds of fusion.

In cases where there is a fusion between tracks that affects more than one operator (for example, if a region is supervised by two operators or there is a fusion between two regions supervised by different operators) the slider levels for the fusion are taken so that the area of operator-assisted fusion is maximised. If the fusion is operator-assisted, only one operator receives it.

At the bottom there are three options regarding the timing of datafuse making fusion decisions.

The first option tells Datafuse how quickly it can make it's initial fusion decision. See section 2.6.)

The second option tells Datafuse when it may change its opinion about an earlier fusion decision. (See section 2.7.)

The third option gives a measure of how quickly Datafuse discards old track associations. (See section 2.8.) If this field is left blank, all track associations are weighted equally.

2.6 Initial track fusion decision

When datafuse is sufficiently confident of a fusion to make a decision about it (See Section 2.4) there are still issues to be considered about whether the fusion should be made or not.

2.6.1 Initial decision delay

In the Configure Fusion Automation window (Section 2.5) the "number of track points before the initial decision is made" may be set. In order to do any fusion there must be at least this many points on the shortest track in the group. This option only disallows any fusion before this criteria is met. Datafuse may still decide later that tracks are associated: just not earlier.

2.6.2 Operator Delay Decision

See Section 2.12.4. Datafuse may not make an association involving a track with a delay on it until the delay period has expired.

2.6.3 Joining track to group

If a track may be joined to an already existing group without reducing the confidence below the operator-assist confidence threshold, it is allowed.

2.6.4 Track is locked

If a track is locked no automatic or operator fusions (other than the fusion process that locked it) may involve it.

2.7 Changing fusion decisions

If the fusion confidence of a group drops sufficiently, a change in the fusion decision may be required. In the Configure Fusion Automation window (Section 2.5) the "Change opinion after drop in confidence" option gives the confidence levels for re-evaluation of fusion decisions.

If tracks have been automatically fused, this is how much the score must drop below the automatic fusion confidence for datafuse to queue the associations on the operator-assist queue. It does not undo the associations itself.

If tracks have been fused by operators with the approval of datafuse (ie. if the overall confidence of the fusion is better than the operator-assist threshold) and the score later drops below the operator assist threshold by the set amount, the associations are requeued on the operator-assist queue.

If tracks have been fused by operators despite a datafuse confidence below the operator-assist threshold, datafuse will not revise the operator decision unless the datafuse confidence later rises above the automatic fusion threshold. After that, the association is treated as an operator assisted association with datafuse approval.

2.8 Weighting Old Associations

Confidence for the association of a pair of tracks is accumulated over the length of the track. The weighting of how much each track point contributes to the overall confidence of the association is a complex issue. If all individual track point associations are equally valid there is no reason to give more weight to any of them in the overall confidence.

Practically, however, it is possible that a change in conditions will reduce the validity of a previous association confidence. In this case it is better to weight more recent points more heavily.

The weighting of a confidence is done by having a time-constant that indicates how long it takes for 36% of the confidence to degrade.

By default, the time-constant is fairly large but circumstances may cause it to be temporarily changed. (See Section 2.9.1.) The first few associations contribute equally towards the confidence before the time-constant comes into effect.

2.9 CR - Datafuse Interface

There are several interactions between the CR System and Datafuse.

Datafuse must request a list of ionospheric modes for a given region for it to do associations of tracks in that region.

It receives a list of the form (Mode Name, Mode Strength, Spread Index, TIP)

The Mode Strength provides a measure of how likely the mode will produce a track in the presence of a target. If datafuse observes no multimoding for a track, the track will be placed in the strongest mode on the list.

The spread index provides a measure of the doppler spread that should be observed on a particular track point that exists in this mode. This aids Datafuse in making (otherwise ambiguous) mode assignments.

TIP = temporal integrity parameter.

2.9.1 Temporal Integrity Parameter (TIP)

The TIP indicates how far into the past that the current CR for this mode should be used. It is a measure of the previous reliability of this mode: if it is zero, the previous CR of this mode is considered just as good as the current. If it is larger, the implication is that previous CR is not accurate.

If the TIP is greater than zero, all of the CR on previous track points from (current time - TIP) in this mode are redone. It is not practical to recalculate track associations, so the association time-constant (see Section 2.8) is reduced proportionally to the inverse of the TIP.

2.9.2 CR Transforms

Datafuse requests the CR for a given track in a given mode:

(ground range, ground azimuth) = slant-to-ground(slant range, slant azimuth, mode name, region number)

or

(slant range, slant azimuth) = ground-to-slant(ground range, ground azimuth, mode name, region number)

2.9.3 CR feedback

CR feedback allows the CR to adjust the height and tilt of ionospheric layers in response to knowledge of ionospheric conditions gained from the fusion of tracks. It is turned on and off by the ECA. (See Section 2.1.1 and Section 2.12.2.)

There are two sources for Coordinate Registration feedback from Datafuse.

Associations between OTHR tracks and ground-registered data (such as ADS data, beacons, Microwave Radar tracks or airplanes) give an absolute frame of reference. CR feedback is of the form:

mode name, time, ground range, ground azimuth, relative range, relative azimuth, range accuracy, azimuth accuracy
where

- mode name = mode we are providing feedback for
- time = time of this association
- ground range and azimuth = OTHR track ground range and azimuth. These are provided so that the CR knows where the suggested change in ionospheric conditions is for.
- relative range = OTHR track ground range / registered track ground range

- relative azimuth = OTHR track ground az - registered track ground az
- range accuracy = combined standard deviation in range of both tracks.
- azimuth accuracy = combined standard deviation in az of both tracks

Relative range can be interpreted as how much higher we think that the ionospheric layer combination forming this mode should be. The range accuracy tells how certain we are of the amount of change required.

Relative azimuth is somewhat more tricky. An azimuth not where it is expected implies either inaccuracy in measurement or ionospheric tilt.

Associations between multimoding OTHR tracks can still give a valuable source of CR in how separated the ionospheric layers are. This is in a relative frame of reference. CR feedback is of the form:

(model name, mode2 name, time, approx. ground range, approx. ground az, relative range separation, range accuracy)
where

- mode 1 and mode 2 = two ionospheric modes that are multimoding
- time = time of this association
- approx. ground range and az = approximate ground range and azimuth. We are not necessarily sure of the exact ground range and azimuth because the mere fact that tracks are multimoding does not automatically register them to ground. These parameters are provided so that the CR knows where the suggested change in ionospheric conditions is for. As ionospheric conditions are usually of a large scale, the exact ground coordinates of the proposed change are not really necessary.
- relative range separation = expected difference in range between two tracks - actual difference in range
- range accuracy = combined standard deviation in range of both tracks

Relative range separation can be interpreted as how much further apart we think the ionospheric layers forming the two modes should be.

2.10 Tracker - Datafuse Interface

Datafuse reads track data from the Stage 3 common after it gets put there by the tracker. While there is potential for increased tracker fusion interaction - such as a doing multimoding at the same time as tracking - this is not currently planned due to the relative difficulty level.

2.10.1 Tracker feedback

The most immediate and simple tracker feedback is for Datafuse to start tentative tracks where it considers there should be targets but there are none visible. Obviously there is the danger of self-fulfilling what Datafuse wishes to see. The tentative tracks started up are recorded by Datafuse, and must have a greater than usual track confidence before being a candidate for fusion.

Datafuse will start up a tentative track only if the ionospheric mode strength and (if known) target size indicate that that there should be a track there, but there isn't.

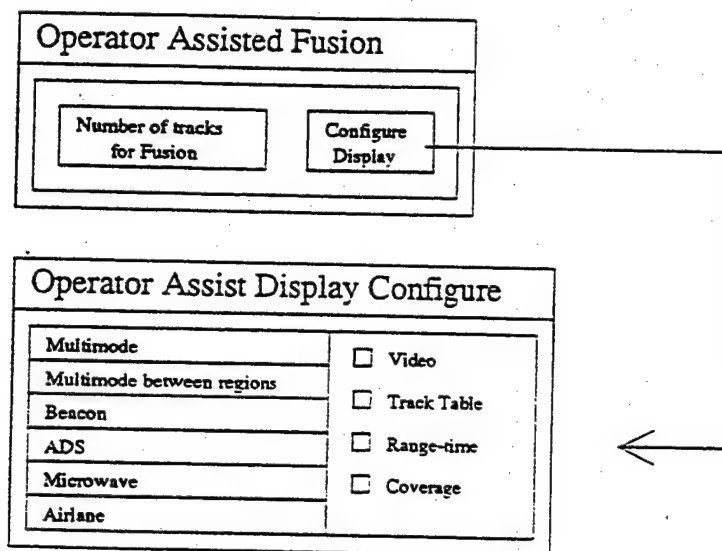


Figure 2: Operator Assisted Fusion Window.

Tentative tracks will be started only once for each expected target position. If the tentative track fails, Datafuse will not try again with that target in that assumed mode.

Datafuse will start tentative tracks when there is an ADS track, Microwave Track, or beacon without a corresponding OTHR track, or when there is multimoding expected but not visible. In the latter case, starting the tentative track is intended to help resolve track-mode ambiguities.

2.11 Operator Assisted Fusion Window

In cases where the fusion algorithm is insufficiently confident of an association to automatically fuse the tracks, but is sufficiently confident that the association is worth considering, it allows operator assisted fusion. (This confidence level is selected in the Fusion Automation window.)

Operator Assisted fusion is handled by the Operator Assisted Fusion Window (see Figure 2).

Operator assist button The number is the current number of potential track associations ready for consideration by the operator. Clicking on the button will bring up the Fuse Tracks Window (See Section 2.12) with the tracks for consideration in it. It will also change a set of the displays to show the tracks in question, and "hook" the tracks on those displays.

Configure Display Window The set of displays changed is configurable according to operator preference, using the configure displays button to bring up the "configure operator assist displays window".

On the left-hand side, the type of fusion is selected. For that kind of operator assisted fusion, the operator may select from the displays on the right which ones get modified to display the tracks being fused.

For multimoding tracks in the same region the default will be the video window and track table. For multimoding between two regions, fusion with beacons, ADS, airplanes and Microwave radar the default will be the coverage window and track table.

Fuse Tracks									
Group Number				Recalculate		Clear Selection			
Overall Confidence		Overall Mode Confidence		Delay decision		for		2 track points	
Track		Track							
Confidence		Confidence							
Mode		Edit		Mode		Edit			
Confidence		Confidence							
Target Type					Tell				
<input type="checkbox"/> Fixed Wind <input type="checkbox"/> Rotory Wing <input type="checkbox"/> Transponder <input type="checkbox"/> Surface Strength <input type="text"/>					<input type="checkbox"/> 2CRU <input type="checkbox"/> CR feedback <input type="checkbox"/> MIC <input type="checkbox"/> Tracker Feedback <input type="checkbox"/> Do not suppress External Number <input type="text"/>				
OK					CANCEL				

Figure 3: Fuse Tracks Window.

Once the operator has closed the Fuse Tracks Window the changed displays will revert to showing what they were showing prior to the operator assisted fusion.

2.12 Fuse Tracks Window

The Fuse Tracks Window (see Figure 3) will replace the "Complex Validate" window in the current displays, and will provide enhanced functionality. It will be able to be selected as an option in the edit function, and will have a button on the toolbar. This window is not a dialog/popup window that must be dealt with before returning control to other windows. The operator should be able to work with other windows independant of the existance of the Fuse Tracks Window, except to bring up another Fuse Tracks Window. The window will disappear when the "OK" or "CANCEL" button is pressed.

At least one track must be hooked before the window may be called. If any of the tracks hooked (or selected later from within the window) is already being dealt with by an operator (is "locked" by the operator) an error message will be printed to the message window and that track will not be included in the list of tracks in the window.

Tracks become "locked" when they are entered into a Fuse Tracks Window or queued

The upper section of the window has a list of tracks and their association and mode characteristics. The lower sections are the same as in the current "complex validate" window target type and who-to-tell.

2.12.1 Target type

The target type section is left in the fusion window because it is relevant to a track group (tracks are only grouped if they are caused by the same target) and because identification fusion may later assist in the identification of target types.

2.12.2 Who to tell

The who-to-tell section is left in because it is also still relevant to a group of tracks and because fusion of OTH with microwave tracks may involve telling different agencies than just multimoding OTH tracks. It is not yet clear what to set the defaults to.

In addition this section indicates whether to give CR feedback and tracker feedback or not. The DTO's by default are not allowed to give either of these (see Section 2.5), but the ECA, using the same display, may set either of these options.

Active CR feedback will cause datafuse to send requests to the CR system to change heights of ionospheric layers, or distance between ionospheric layers based on the associations made. Association of OTH tracks with beacons, Microwave radar tracks, ADS tracks, airplanes etc. will give absolute coordinates of tracks, and will produce feedback to the CR that will be able to change ionospheric layer heights in the CR matrix. Multi-mode associations do not give absolute coordinates of tracks, but will allow us to adjust the distance between the ionospheric layers.

2.12.3 Tracks and Associations

Any tracks that were hooked when the fuse-track window was called will be listed in the table. In addition, if any of the tracks hooked represents a "track family", all of the tracks that are members of the track family will be listed in the table as well. These tracks - normally invisible in the range-time, video and coverage windows (see Section 3.1 and Section 3.2) - are set to be visible on the display station that the fuse-tracks window was called on, so long as the fuse-track window is active. This is so that the operator can see what tracks went into making the track family.

It will be possible to hook and list non-OTH tracks and have them come up on the table. (For example, microwave or ADS tracks.) These tracks will have blank mode fields.

Datafuse will evaluate the list of tracks as a hypothesis group, and return the confidence that the tracks are associated (under the track number), the mode that the track is in for this to occur, and the confidence that the mode chosen is the correct mode for that track (under the mode). It will also output the overall confidence that the selected tracks represent a group, and the overall confidence that the modes selected for these tracks are correct, given that the tracks are associated. A group number will be generated if one does not already exist.

2.12.4 Delay Decision

If the operator is not certain of the fusion and wants to wait a bit to see if after a few more points the associations are more firm, they may click on the "Delay decision" button. The amount of time that the delay is for may be varied in the window just next to the button. This indicates the number of points all tracks must get before the association may be requested. In other words, the effect of the "Delay decision" button is to disallow any fusion decision until the required number of track points has appeared on each track. The fusion is then re-evaluated.

2.12.5 Editing the fusion decision

At this point the operator may edit the list of tracks. They may add or delete tracks from the table, and change or delete modes.

Editing any of the tracks or modes will clear the confidences returned by datafuse. If the operator wishes to see what datafuse thinks of their selected group of tracks, they may press the "recalculate" button, and datafuse will recalculate its confidences based on the new list of tracks and modes. Datafuse will not change any modes that are still defined at this stage, but will interpret any blank mode field as a request to select the best mode for the track.

At this time, as new tracks may have been entered into the track table, it also checks that all of the tracks in the table are locked by it. Any tracks that are not locked become locked by it. If a track is locked by another process it is removed from the track table and an error message is reported to the Message Window.

Adding tracks is done by selecting a track field and typing in the track number.

Deleting any track or mode may be done by selecting the field and clicking on the "clear selection" button. Deleting multiple tracks and modes may be done by right-clicking on the fields (in the same way as multiple tracks are "hooked") and clicking on "clear selection".

A different mode for a track may be selected by clicking on "edit" next to the current mode. A list of currently available modes will appear, with the confidence that the track is in that mode alongside. The selected mode will appear in the mode window. If the operator chooses a mode that another track in the group has, the other track's mode will be deleted.

When the operator is finished editing they press the OK button and the changes to the associations are made. As above, the process checks that all tracks in the table are locked by it. If any of the tracks are locked by another process, it is removed from the track table, an error message is reported to the Message Window, and the Fuse Tracks Window does not close.

If all tracks are deleted, no changes to the current associations are made. Otherwise, any tracks that get removed from that group are tagged so that datafuse may not rejoin them to that group. Regardless of the datafuse confidences any tracks that the operator has fused will remain fused.

There are some fusion possibilities that are not available to the DT operator. These are fusions that give absolute ground truth as a result of association. (For example, fusion with airplanes, microwave radar tracks, ADS data, beacons, etc.) The reason for this is that the associations give feedback to the CR system, and only the ECA operator may make decisions that give CR feedback.

3 Other Issues

3.1 Invisible Tracks

It is desirable that there be an option to make tracks invisible without breaking or killing them. The tracks continue but are not shown to the operator. This would be an option in the toolkit and the edit function.

It would be useful for the fusion process to hide multi-moding tracks and for the forthcoming track prioritisation algorithm (being developed by Chris Crouch) to suppress low-priority tracks.

It would not be difficult to implement as there are already mechanisms in place for suppressing classes of tracks such as stationary tracks and tracks in particular modes.

Mode
<input type="checkbox"/> Unvalidated
<input type="checkbox"/> Confirmed, Fused Tracks
<input type="checkbox"/> Tracks forming fused track

Figure 4: Track Overlay Window (extract).

3.2 Track Overlay Window

In the current system it is possible to configure which tracks are visible on the operator displays using the Track Overlay Window called from the options menu. The section labelled "Mode" allows various modes of tracks to be suppressed. This part of the Track Overlay will be changed as shown in Figure 4.

If "Unvalidated" is set, unvalidated tracks will be shown on displays. This is set by default.

If "Confirmed, Fused Tracks" is set, confirmed and fused tracks will be shown on displays. This is set by default.

If "Tracks forming fused track" is set, tracks that are members of track families are shown on displays. This is unset by default.

3.3 Display of Fused Tracks

Fused tracks in ground co-ordinates do not actually exist as far as the tracker is concerned, but we still wish to display them. We also may wish to display the fused track converted via the dominant mode in OTHR coordinates.

Likewise, we may wish to display Microwave or ADS tracks on a ground coordinate display without interacting with the tracker. These tracks may still be fused with OTH tracks or each other, so it should still be possible to "hook" them.

3.4 Other Track sources

The fully integrated data fusion system requires access to track sources other than OTH - for example, Microwave radar data, ADS data, airplane positions and beacon data. Some of the data is already accessible on the system (airplane positions and beacon data) but it is currently not clear whether the others will be integrated. It is possible, for example, that Microwave radar tracks will not be given to JFAS in real time.

If this is the case, fusion of external source tracks becomes more complicated and less accurate. If data fusion is still required with the external source tracks, the fusion will have to be done externally. Datafuse requires a large set of OTHR specific data in order to function properly, of which only a subset is accessible externally. It would operate with considerably less accuracy if the only data from JFAS was OTHR tracks. It would be impossible for it to do multi-moding without CR information.

A proposed system has one version of datafuse running within JFAS and assigning modes to tracks, bringing the tracks to ground. A second version, using only the OTHR ground tracks and external tracks, would do fusion purely in ground co-ordinates. This system has the disadvantage of not being able to use the maximum set of track information. For example, ambiguity in multimoding could be aided considerably by having Microwave tracks in the same area.

If the external sources of data are integrated, there are still issues that have not yet been explored. It is obvious that the tracks will be displayed, but not so obvious who will be able to view them (ECA?) or how they will be distinguished from other tracks in displays (different colours?) and track numbers (letter before the track number?).

4 Implementation of Changes

The above features will require large changes to the current displays and stage 3 architecture. It is proposed that the implementation be done in several stages so that the changes can be made gradually.

Almost all of the below changes made to get to an intermediate version should be useful in the complete version.

4.1 Version 1.0

4.1.1 Functionality

Version 1.0 is designed to enhance the multi-mode recognition capability only.

The only display change will be the introduction of the Automatic Fusion Window in the DTS display (See section 2.1). This will have only one slider on it: the operator assisted fusion option will not be available. The slider will denote the confidence at which automatic fusion will occur.

The configuration option of the Automatic Fusion Window will have reduced functionality. Only the bottom two options will be active: the "Associate after X points on track" and "Change opinion after drop in confidence of X %". The Automatic Fusion slider value and these two options are valid over the full coverage.

Datafuse will do complex validates involving only unconfirmed tracks.

It may change its opinion on an association if the confidence drops sufficiently. If it does so, it may undo any complex validate that it created, but not one that the operator created. If it undoes a complex validate it sets the tracks back to "unconfirmed".

Datafuse will use its own backup CR system rather than the new CR system. It will use operator generated ionospheric heights to calculate modes for the fusion.

4.1.2 Architecture

Figure 5 shows the proposed architecture for the first step of datafuse implementation.

4.1.3 Sequence of events

- Central Update informs Datafuse when Stage 3 has finished writing the stage 3 database (Circled 1 in 5.) If Datafuse has not finished on its current list of track associations it complains, and dumps the remaining unprocessed track points.
- Get tracks from the stage 3 common (2)
- Get region details of the tracks from central update. (1)
- Get the operator-set ionospheric heights and create ionospheric modes from them based on rules of thumb. (1) If the operator ionospheric heights indicate no multimoding is taking place, datafuse selects two preser ionospheric heights and creates three modes from them.
- Make associations between tracks.

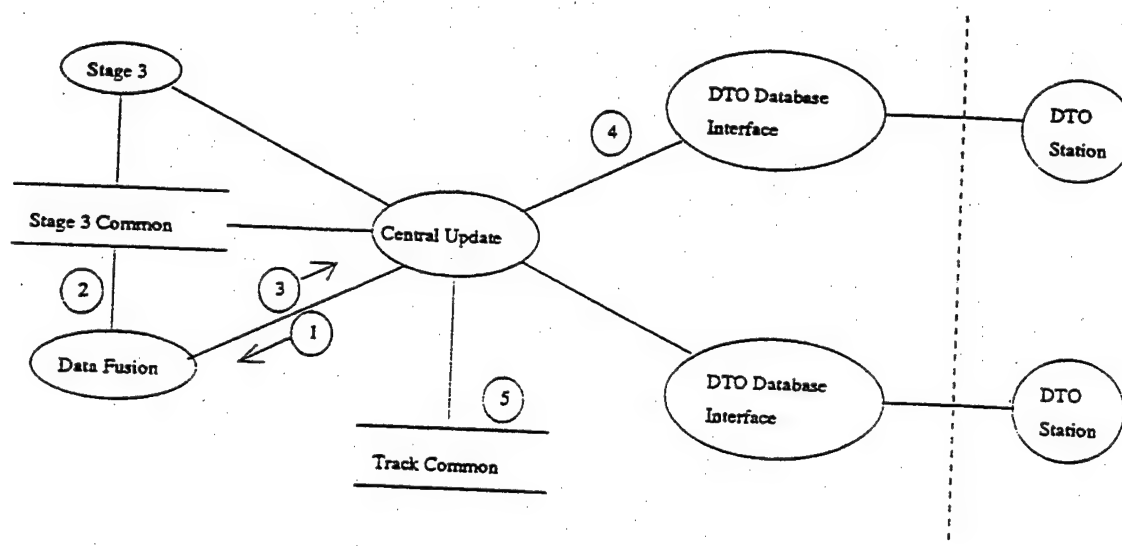


Figure 5: ERP stage 3 architecture with datafuse.

- Read the track common to determine which tracks are unvalidated and thus able to be fused. (1), (5)
- Get current automatic fusion confidence, number of track points for first association, confidence drop level for opinion change from the Automatic Fusion Window via Central Update. (1), (4)
- Determine which tracks to complex validate. Send "complex validate" instructions of tracks that it considers associated to central update. (3)
- Determine which previous decisions have dropped in confidence sufficiently to undo them. Send instructions to unvalidate the relevant tracks to central update. (3)

5 Discussion

Who-to-tell section of the fuse-tracks window: this should change with the tracks being associated. eg. if a microwave track is one of those being associated we should have a different default "tell" than if all the associated tracks are multimoding OTH.

Should DT operators be able to give CR feedback? If DT operators are allowed to do multimoding this can still give relative feedback of the form "These layers are too close together by 10%".

Perhaps all operators could do fusion that gives feedback to CR, but the ECA is automatically consulted before the fusion is accepted?

Tracker feedback: using knowledge of multimoding tracks to improve their tracking.

Where should datafuse go? Will JFAS get microwave data from the sadoc or will we have to send OTH data there and do the fusion there? A lot of the above assumes that JFAS will get microwave data. Is there any plan to integrate the Microwave stuff into the displays?

DRAFT

Recent improvements to datafuse

Andrew J. Shellshear

October 12, 1995

1 Treatment of regions

Region data provides valuable information in the association of tracks that has previously been ignored. The resolution of track parameters depends on radar operating frequency. WRF and other information that varies from region to region. Previously resolution of track parameters has been fixed for all tracks at compile-time.

The new treatment of regions has other benefits. Ionospheric modes vary with operating frequency and hence with region. Consequently ionospheric modes are now stored regionally and interaction between tracks of different regions can take this into account.

This has had profound affects on the way datafuse does it's associations. Each region can be treated as analagous to a separate radar with separate operating conditions but a common transmitter and receiver site. Associations across regions become a simple case of multi-radar association.

So, association of a real radar track with a hypothetical radar track (produced by calculating a hypothetical ground position of a radar track and generating a possible radar track based on an ionospheric mode) has an additional complication. The hypothetical track must belong to a particular region (which has it's own ionospheric conditions). Hence there will be a hypothetical track for each ionospheric mode in each region, an expansion in number of hypothetical tracks of the order of num_regions. Fortunately in most cases, the hypothetical track created assuming the ionospheric conditions of a particular region will produce tracks that are not at all close to the region. The partition grid has a "region proximity" list which lists nearby regions to each partition. If the given region is not nearby, the hypothetical track is ignored.

2 Interpolation and extrapolation of points

Unless tracks being compared are from the same region it is unlikely that the track points are at the same time-stamps. Previously we have always done interpolation of tracks to enable comparisons at the same time-stamp, but this often requires waiting until there are more points on the tracks before doing associations. Associations would always be out-of-date by at least one track point.

Recently the ability to extrapolate has been added. Extrapolation is only done if the difference between the time-stamps of the two track points is small.

In most cases, this has meant that there is no wait necessary before associations, and the delay in output of associations is limited by only the speed of datafuse.

3 New association scoring system

Previously we have used a simple sum to get overall association scores for the length of a track. The overall score is `sum_of_scores` divided by `number_of_comparisons`. While this is theoretically satisfying - in the long run, no point on the track is more important than any other point, so it gives the clearest "overall picture" - in an operational sense, the most interesting and important data is the most recent.

Both points of view have merit. The second one initially seemed like a placebo for the operators - but in a practical sense, weighting recent track points more strongly than old ones is a good idea. The tracker may have merged two tracks; for example, or the CR may have been incorrect for a while and then suddenly fixed, creating much better or worse associations. The point is, if mistakes have been made, weighted track scores will respond much more quickly to corrective changes (and to mistakes too, of course - which is not always a bad thing.)

So it makes sense to store both scoring systems, the previous

`total_new = score + total_old`

and the new

`total_new = momentum * total_old + (1 - momentum) * score`

where momentum describes how much previous points contribute to the score, and can be expressed as

`momentum = exp(-1 * delta_time / time_constant)`

where `delta_time` is the amount of time since the last update of the score, and `time_constant` is the amount of time for 63% degradation of the previous score.

Note that because we do twice as many comparisons as we expect (because of the track time-offset possibility, we do two comparisons for every pair of track points - even when the track points are simultaneous. This is for compatibility.

—time→

—0—x—0—x— track1

—x—0—x—0— track2

x = track point

0 = extrapolated or interpolated point for association.

There are two points on each track but four comparisons made. Thus, the actual time-constant used is half the stated time constant, to get the stated time-constant's effect.

4 CR feedback

To give the possibility of CR feedback, for each track association the difference between actual and hypothesis positions are accumulated over time. This will give long-term trends in the accuracy of the CR system.

ALICE: A Fuzzy Logic Approach to Automatic Multi-Sensor Data Fusion

W.T. Johnson and D.J. Kewley,
High Frequency Radar Division,
Defence Science & Technology Organisation, Department of Defence, Australia

Abstract: *The general multi-sensor data fusion problem is to take data from multiple sensors and to usefully combine it in real time, to provide information that was not available from any individual sensor. A particular example of interest to the Department of Defence is discussed; the fusion of data from a Microwave Radar (MR), and an Over-The-Horizon-Radar (OTHR) located at the Jindalee Facility Alice Springs (JFAS). The new fuzzy-logic data fusion engine ALICE (Always Logical In Connecting Events) is discussed together with the results of the fusion of actual MR and OTHR data and the results are compared to ground-truth data. This is the first widely reported time[1] that a fuzzy-logic method has successfully combined sensor track data and this is demonstrated using OTHR and MR data. The same method abstracts and classifies, track data and this is demonstrated using some Automatic Dependent Surveillance (ADS) data.*

Keywords: *fuzzy logic; data fusion; combination; abstraction; radar.*

1 INTRODUCTION

Data fusion, in its most simplistic form, is the combination of information from two or more sensors to determine a level of knowledge greater than that obtained from the simple sum of knowledge obtained from the individual sensors alone. A particular example of interest to the Department of Defence is the fusion of data from a line-of-sight Microwave Radar (MR), and an Over-The-Horizon-Radar (OTHR) located at the Jindalee Facility Alice Springs (JFAS). Many approaches to the multi-sensor data fusion problem are possible, and are under active investigation. These methods include the multi-hypothesis-method (MHM), the artificial-neural-network-method (ANNM) and the fuzzy-logic-method (FLM). The MHM[2,3] has been used to automatically solve the multi-mode problem for OTHR and work is currently underway to extend this approach to other sensors. Through a DSTO sponsored project, an ANNM has been proposed[4] to solve the multi-mode problem. Both of these methods amount to connecting events spatially but not temporally. In this paper just one approach is discussed in detail, namely the method of fuzzy logic. The ALICE system (Always Logical In Connecting Events) is based on previous work at HFRD; philosophically it draws on Kewley[5] and in its implementation on, Kewley et al[1,6,7], which was capable of connecting events spatially and temporally. Implicit in all of these data fusion methods, is the assumption that the OTHR data have already been formed into tracks[8,9].

1.1 Current Operation

Currently, the two sensors, the OTHR and the MR, are operated separately. Each produces its own automatic tracks, which are validated by operators and are subsequently manually combined. However, there are unique characteristics of OTHRs which make this method of combination inefficient and these will be identified in the following section.

1.2 Traditional Operation Of The JFAS OTHR

In broad terms, the JFAS OTHR operates[10] by transmitting radio waves, slanted skyward above the horizon, which are then refracted through the ionosphere and back toward the earth surface, far over the horizon from the transmitter. The radio waves scatter from earth surface and any other object, say an aircraft, that happens to be present and return back to a receiver via the ionosphere for a second time. Doppler processing detects moving objects. The length of the radio wave propagation path from the aircraft is called the slant-range, and since other parameters, such as speed, are measured with respect to this path it is known as *slant-space*, and it is the natural coordinate system of an OTHR. Given the ionospheric height midway between the receiver and the aircraft slant-space parameters can be transformed to standard geographic coordinates of latitude and longitude. To a first approximation the refraction can be treated as a simple reflection and the corresponding height is called the virtual ionospheric height.

Errors can occur in this conversion from slant to geographic coordinates. One from mistakenly choosing the wrong ionospheric layer, known as *miss-mode*. Another from using the wrong virtual height when using the correct layer, known as *height-assignment error*, and finally when multiple tracks occur due to multiple layers, which is called *multi-mode propagation*. Multi-mode propagation, and miss-mode and height-assignment errors are identified by operators using various tools and displays, but errors still occur in practice.

1.3 Multi-Sensor Data Fusion

There is scope for improvement by utilising other sensors; that is to say by data fusion, although clearly it is limited to the regions of overlapping coverage. However the benefits of data fusion are not limited to assistance in the OTHR coordinate registration (CR) process. A clear unified air traffic picture may also be gained, with a single track per aircraft, by taking tracks from each sensor and combining them together. So called negative information may also be utilised. For example,

although an OTHR does not normally yield an altitude estimate, within the nominal MR coverage, an OTHR track without a corresponding MR track implies that the aircraft is flying very low beneath the horizon of the line-of-sight MR. Implicit in a complete data fusion scheme is the process of classification and abstraction; the extraction of essential features from the mire of all the information simply lumped together. In military terms this is called 'Situation Assessment'. For instance, the aircraft #1 is 'circling', and the aircraft #2 has 'manoeuvred off the airframe' are clearly the essential characterisations or encapsulations of a series of the usual reports of track-number, time, latitude, longitude, heading, speed and altitude (TT-LL-HSA).

The architecture of data fusion[11,12,13,14,5] is commonly divided into the *four US JDL levels of data fusion*: Level 1 *Object Refinement*, Level 2 *Situation Assessment*, Level 3 *Threat Assessment* and Level 4 *Product Improvement & Sensor Suite Management*, and additionally into *four stages of data fusion*, *alignment*, *correlation*, *association* and *combination* which are applied in each of the bottom three levels. Therefore there are two processes at work, firstly, the tendency toward combination of data at that level and secondly, the tendency to abstract information up to a higher level. In this paper the combination of MR and OTHR track data constitutes Level 1 Object Refinement, while the goal of a reasoning with tracks represents Level 2 Situation Assessment and it will be shown that the method presented here allows both.

The first step in multi-sensor data fusion is to understand the sensors. A pulse MR for instance may have accurate position but noisy speed and heading estimates. In contrast a Frequency-Modulated Continuous-Wave (FMCW) OTHR may produce accurate radial-speed or doppler, but noisy position estimates. However, over and above these considerations the process of coordinate registration, as outlined previously, is *critical* for OTHR and must be carefully taken into account. Additionally the sampling-rates will vary for OTHR depending on operating conditions and tasking constraints, whereas for MR it will be fixed at a relatively high rate.

1.4 The Fuzzy Logic Approach

Clearly, the fusion of OTHR & MR data is a complicated problem involving complex sensors and human interaction. It is proposed to solve this problem and manage the complexity by using Fuzzy Logic. In fact, one goal is to emulate a human *combining and reasoning* with tracks from multiple sensors of different source types. Conventional trackers can combine, but cannot reason about the behaviour of the entities being tracked with respect to each other; humans do both. Therefore, rather than physically modelling the problem, a *systems approach* is used and the *human reasoning chain* is modelled by using a *fuzzy rule-based system*. The generalisations available by using a fuzzy rule-base system are due to the wide scope of the rules, including coping with non-numeric information. For example, the additional information from a visual sighting 'they were both small grey planes' does not fit well into a tracker state-space vector, but by contrast, is easily incorporated as a fuzzy rule. Thus the overall goal can be clearly seen; the attempt is to enhance the system performance of a suite of sensors, of different source types, by *adding a higher level stage*. Additional fuzzy rules could be developed in cooperation with the Australian Defence Force (ADF) to allow them to contribute their knowledge and experience directly.

2 CLASSES OF CONNECTING EVENTS

In the ALICE system, a number of events, such as a report containing TT-LL-HSA, may be connected in certain classes. Consider two TT-LL-HSA events or reports, pertaining to aircraft. A human doing some 'back of an envelope calculations' would connect them based on how well the *expected*, or current, speed and heading match the speed and heading *estimated* by using the ground-range, bearing and time-difference between the reports. If the two events almost exactly match or connect by this calculation then it is dead-reckoned to be on-course, or simply DEAD-ON-COURSE. If the match is a little worse it is called NEW-ON-COURSE. Alternatively these could be written as ON-COURSE(DEAD) and ON-COURSE(NEW) to show how the terms in parenthesis are modifiers on a root class. This class takes into account the worst errors that can be expected from the sensors with the caveat in the case of OTHR of reasonable CR. If events do not match in these classes and one of the sensors is an OTHR, the effect of poor CR information can be tested.

Given the virtual ionospheric height used for the track, the mode can be classified as E-LAYER, F-LAYER, F2-LAYER, E-F-MIXED-LAYER, F-F2-MIXED-LAYER¹ and then the alternatives are considered. To a first approximation an E-layer is about 100km, the F-layer 220km, the F2 is in the range 300 to 350km and the mixed layers are roughly halfway between the nominal values. For instance, suppose the virtual ionospheric height given is 157km. This is roughly the equivalent height of an E-layer transmission out, and an F-layer back. Therefore, in this example, the alternatives are to bias the current height of 157km by -50km to account for a possible E-layer, +50km to roughly approximate a pure F-layer, +100km to approximate an F-F2-mixed-layer, and finally +150→+200km for a pure F2 layer. Now the OTHR data in slant-space are brought to ground with these biased virtual ionospheric height values. After taking into account the likely spread of values about the nominal values, the possibility for a connection between events can be tested using the similar classifications to the above, such as, ON-COURSE(WITH JFAS-BIAS +50Km), or ON-COURSE(WITH JFAS-BIAS [E-F-MIXED-LAYER] TO [F-LAYER]) where

¹ Alternatively a better CR system might deliver this information directly.

again the terms in parenthesis are modifiers on the class ON-COURSE.



Finally, tests are made for a number of manoeuvring behaviours. The first series are as modifiers to the On-Course classes, such as ON-COURSE(MANOEUVRED) to indicate a match on speed but a change of heading, or ON-COURSE(ACCELERATED) and ON-COURSE(DECELERATED) to show matches in heading but speed changes. Proximity is limited in this type of modifier to keep the number of matches tractable. The class DRASTIC-MANOEUVRE, with the modifiers (FORWARD), (BACK-TRACK) and (CIRCLING) are the final series of matches considered. The first two modifiers are largely intended to connect events which appear as improbable drastic large-scale manoeuvres because they are reported sparsely in time, while the third connects events associated with circling aircraft.

This discussion of the classes appropriate to the system together with their hierarchy, although brief, is actually a complete description of ALICE. For each class a single nominal value has been identified, together with some sort of error value which describes the support or extent of the class. What is required now is a mathematical representation of the various classes and also allows manipulation and comparison of the various quantities.

3 A FUZZY MODEL OF IMPRECISION, CLASSIFICATION & ABSTRACTION

A traditional way of handling classes is set-theory or its isomorphic counterpart, logic. However, it is also desirable to weigh a set-member, at the nominal value of a class more heavily than one near the extreme. Clearly, this is an application well disposed to the use of fuzzy-sets; Kewley et al[1,6,7] have prototyped this approach in the Prolog language. They calculated a nominal heading value C_{H_s} , which characterises the set, based on the bearing between two reports, together with a representative error value, E_{H_s} . Note that the subscript s refers to estimated values, while the subscript e will refer to expected values. The bounds representing the support of the set, H_s , are taken as $C_{H_s} \pm E_{H_s}$. Thus, as shown in Figure 1, for all aircraft headings h in $[0, 360 \text{ degrees})$, $0 \leq \mu_{H_s}(h) \leq 1$ indicates how close a particular value of h is to the nominal value C_{H_s} . Thus if the set, H_s , is defined on a universal set of headings H , which is finite and countable, and $\mu_{H_s}: H \rightarrow [0, 1]$

$$H_s = \sum_{i=1}^n \mu_i/h_i, \quad \text{where } \mu_i = \mu_{H_s}(h_i).$$



Similarly, another fuzzy set can be formed for H_e , the expected heading. For the first point on a track the error value E_{H_e} would be based on the sensor type and is in this case identical to E_{H_s} . Thereafter, it will be modified according to the precision of the match; this will be defined later. Figure 2 illustrates the possibility that H_s is H_e . Note, that there three alternative rule types, truth-qualifying, certainty-qualifying and possibility-qualifying rules[15]. In the ALICE system possibilities are used. Thus the possibility[16,17,15], that H_s is H_e , is given by

$$\text{Poss}(H_s \text{ is } H_e) = \sup[\min(\mu_{H_s}(h), \mu_{H_e}(h))] \quad \text{for } h \in H_s, H_e.$$

Additionally, other sets for V_s and V_e , the estimated and expected speed, can also be formed. Thus following the standard operators in fuzzy reasoning[16,18,19], the membership-value that the pair of events (e_s, e_e) have in the set ON-COURSE is given by

$$\mu_{\text{ON-COURSE}}(e_i, e_j) = \min[\text{Poss}(H_i \text{ is } H_j), \text{Poss}(V_i \text{ is } V_j)].$$

That is to say if $\mu_{\text{ON-COURSE}} \neq 0$, then $\{H_i \text{ is } H_j\}$ AND $\{V_i \text{ is } V_j\}$, so it is *possible*, but probably not *certain*, that the latest event e_j is the connected with the last event e_i , in a track t_n , in the class of ON-COURSE. All that remains is to choose the associated track t_n based on the best correlating match from the tracks t_n for $n=0 \rightarrow N$, by implementing the new rule hierarchy outlined in section 2. Note that, in direct contrast to previous work[1,6,7], a breadth first search is used for speed, while a restricted depth search is added on some occasions to avoid seduction due to the inevitable delivery of non-time-aligned events across the various sensors, which results in some reordering of connections in the most recent series of events. This new search method allows the tracking of entities through time, whereas previous practice[1,6,7] only dealt with point by point comparisons. This, together with a stronger sense of entity identity sets the current approach apart from Kewley et al[1,6,7].

The errors E_{H_e} & E_{V_e} used in the sets H_e & V_e are calculated from the time-difference Δt_{ij} and ground-range G_{ij} , between the events (e_i, e_j) , and the expected error in location E_{L_e} for the class ON-COURSE(DEAD); where E_{L_e} has been taken from the track. For the sets H_e & V_e the errors E_{H_e} & E_{V_e} are also taken from the track. For the first point on the track the expected errors are identical to the sensor-type errors. In the case of ON-COURSE(NEW) the *Max* of the sensor-type errors are used for E_{H_e} , E_{V_e} & E_{L_e} and the sets H_e & V_e are recalculated accordingly. The extent of support of the heading set for the class DRASTIC-MANOEUVRE(FORWARD), H_{mf} , which is still centred on C_{H_e} , fills the universal set H , ie. $E_{H_{mf}} = 180$. Likewise, for the velocity V_{mf} the lower side error extends to zero, but the upper side is limited by a maximum manoeuvre speed, v_{nm} , to place a reasonable constraint on speed, following previous work[6,7], but note their single manoeuvre set was different. The back-tracking set membership function is calculated as $\mu_{H_{mb}} = 1 - \mu_{H_{mf}}$. The accelerate and decelerate sets are based on lower and upper halves of V_{mf} while the circling velocity set V_{mc} is given by $\mu_{V_{mc}} = 1 - \mu_{V_{mf}}$ for $v_i < v_{nm}$, and $\mu_{V_{mb}} = 0$ otherwise.

Now that an association has been produced, with an accompanying reason or class, combination can be attempted. Efficient methods of combination are still under active investigation, but at least for a initial attempt the method first proposed by Kewley et al[1,6,7] may be trialled, although they never successfully demonstrated it on a comprehensive scale. For the ON-COURSE classes, the compromise or fused value is taken at the intersection of sets, h_e & v_e for heading and velocity respectively; see Figure 2. The precision, h_p , or new expected error value, is taken in one of four ways involving the four extents from the compromise value to the four zero membership points, combined with *Min* and *Max* functions, on the lower and upper sides of the compromise value. Currently the *Max(Min(extents))* is used as a mid-range value to update the expected errors, E_{H_e} & E_{V_e} . E_{L_e} is also updated based on these values and h_e , v_e , Δt_{ij} & G_{ij} . However, for any of the manoeuvring classes the reported values are used as the fused value, since the history of the track in these classes of connection must be ignored.

4 EXAMPLES





Figure 3 shows a geographic coverage of the fused, or combined, MR & OTHR data. Four tracks are shown; a JFAS OTHR track, with a MR track in their original positions, together with the fused MR-OTHR from ALICE and a Global Positioning System (GPS) ground-truth-track to indicate the real position of the aircraft. Note that although there is a significant range

difference between the OTHR and MR track, the fused track has contributions from both. In fact, the OTHR track-data is significantly reducing the speed-error in the ALICE fused-track compared to the GPS ground-truth, than the raw MR-track compared to the same GPS ground-truth. This is a successful demonstration of a full Level-1 Data Fusion Object Refinement; alignment, correlation, association, and combination.

Figure 4 shows some sparse data from a single Qantas flight delivered via the international ADS system, together with the airline along which it is flying. The geographic coverage plot shows the aircraft at first, travelling along an airline, then manoeuvring, presumably to avoid a storm or turbulence, and subsequently manoeuvring back onto the airline. ALICE correctly connects these events, and while this is trivially easy, notice what ALICE gave for free. She said ON-COURSE(DEAD), ON-COURSE(DEAD), ONCOURSE(MANOEUVER), ie. On-Same-Speed with a significant heading change, ONCOURSE(MANOEUVER), ON-COURSE(DEAD), hence encapsulating hours of flight into an abstract essence. It is envisaged that by adding a rule to cope with the airline and the general region that the encapsulations ON-COURSE_ON-AIRLINE_IN-SEA-AIR-GAP, and DRASTIC-MANOEUVER(FORWARD)_OFF-AIRLINE_IN-SEA-AIR-GAP will be of significant interest to the ADF. This example is used to demonstrate the potential for ALICE to perform Level-2 Data Fusion Situation Assessment.

5 CONCLUSIONS

Significant developments in the use of fuzzy-logic and sets in the area of multi-sensor data fusion since Kewley et al[1,6,7] have occurred, and have been reported here with some early results. The new ALICE system has added and implemented the concept of root class and modifiers or linguistic hedges, to describe behaviour. Also there is a more effective rule structure and a significantly better search method. This allows entity history to influence association and thus combination. Some specific results are that Level-1 Data Fusion Object Refinement has been demonstrated by combining OTHR & MR track data to produce a unified fused track. Further, a rudimentary categorisation of behaviour into classes has been shown, and there is enormous potential for further development of the ALICE system to achieve the goal of giving Situation Assessment, since it already has the tools to perform spatial and temporal reasoning.

Fuzzy sets are proving to be not only a simple and powerful paradigm for modelling imprecision, handling complex sensors and a complicated system, but more importantly are giving, in a quite natural way, the ability to abstract the essence of a situation from the enormous mire of low-level track report information.

REFERENCES

1. KEWLEY, D.J., JARVIS, B.J., NELSON, M.J. (1992), Integration of Jindalee OTHR with External Ground-Truth Data. *Available to qualified readers*.
2. DALL, I.W., SHELLSHEAR, A.J. (1993), Evaluation of a Model based Data Fusion Algorithm with Multi-mode OTHR Data. *Proceedings of Asilomar Conference on Signals, Systems and Computers*, Pacific Grove CA, pp 273-277.
3. DALL, I.W., SHELLSHEAR, A.J. (1994), Performance of an Improved Model Based Data Fusion Algorithm for OTHR Data. *Proc. Radar 94*, Paris, pp 625-629.
4. ZHU, J., BOGNER, R.E., BOUZERDOUM, A., SOUTHCOTT, M.L. (1994), Application of neural networks to track association in Over The Horizon Radar. *Proc. SPIE Symp. on Optical Eng. in Aerospace Sensing*, 2233, pp 224-235.
5. KEWLEY, D.J. (1993), A Model for Evaluating Data Fusion Systems. *Proceedings of Asilomar Conference on Signals, Systems and Computers*, Pacific Grove CA, pp 273-277.
6. JARVIS, B.J., KEWLEY, D.J. (1994), Automatic Multi-Sensor Data Fusion Processing Using Fuzzy Logic. *Proc. AIM-TEC 94*, Adelaide South Australia, pp 26-29.
7. JARVIS, B.J., KEWLEY, D.J. (1994), Fuzzy Logic Application to Multi-Sensor Data Fusion Processing for OTHR Performance Assessment. *Proc. Radar 94*, Paris, pp 105-110.
8. COLEGROVE, B., AYLIFFE, J.K. (1987), An extension of Probabilistic Data Association to include track initiation and termination. *Proc. of the 21st International Electronics Convention and Exhibition, IREECON87*, The IREE Australia, Sydney Australia, pp 853-856.
9. COLEGROVE, B., EDWARDS P.J. (1987), Review of the automatic tracking of targets in clutter. *Proceedings of the 21st International Electronics Convention and Exhibition*, The IREE Australia, Sydney Australia, pp 681-684.
10. SINNOTT, D.H. (1988), Jindalee - DSTO's Over-The-Horizon Radar Project, *Electronics Today International*, Feb.
11. WALTZ, E., LLINAS, J. (1990), *Multisensor Data Fusion*, Artech House Inc. Norwood MA.
12. LUO, R.C., KAY, M.G. (1992), Data Fusion and Sensor Integration: State-of-the-Art 1990s, *Data Fusion in Robotics and Machine Intelligence*; ABIDI, M.A., GONZALEZ, R.C. (editors), Academic Press, San Diego CA.
13. THOMOPOULOS, S.C.A. (1990), Sensor Integration and Data Fusion, *Journal of Robotic Systems*, 7, pp 337-372.
14. PAU, L.F., (1988), Sensor Data Fusion, *Journal of Intelligent and Robotic Systems*, Vol 1, pp 103-116.
15. DUBOIS, D., PRADE, H. (1991), Fuzzy sets in approximate reasoning, Part 1: Inference with possibility distributions & Part 2: Logical approaches, *Fuzzy Sets and Systems*, 40, pp 143-202 and pp 203-244.
16. ZADEH, L.A. (1983), The role of fuzzy logic in the management of uncertainty in expert systems, *Fuzzy Sets and Sys.*, 11, pp 199-227.
17. PRADE, H. (1985), A computational approach to approximate and plausible reasoning with applications to expert system, *IEEE Trans. Pattern and Machine Intelligence*, 7, pp 260-283.
18. KOSKO, B. (1992), *Neural Networks and Fuzzy Systems: a dynamical systems approach to machine intelligence*, Prentice-Hall, Englewood Cliffs, New Jersey.
19. MAMDANI, E.H. (1977), Application of fuzzy logic to approximate reasoning using linguistic synthesis, *IEEE Transactions on Computers*, C-26, pp 1182-1191.

4.4 A COMPARISON BETWEEN US AND AS DATA FUSION APPROACHES

White, Kruger

From: White, Kruger
To: Yssel, Bill
Cc: Torrez, Bill; Radar Agreement Management
Subject: MRTDF and DATAFUSE comparison for MOA
Date: Wednesday, 25 September 1996 5:30PM

Data Fusion PA, file:B7103/2/6

Bill,

Here's the latest version of the comparison report for your input.

Cheers,

Kruger

Fusion of OTHR tracks - A comparison between US and Australian approaches

This draft report compares the methods of fusing OTHR tracks. The two approaches considered at this stage are the US Multiple ROTHr Track Data Fusion (MRTDF) system and the Australian approach, DATAFUSE. MRTDF and DATAFUSE both use the technique of testing multiple propagation hypotheses in order to decide whether tracks are associated. It is intended that other approaches will also be included in the comparison, for example Australia's ALICE (Always Logical in Connecting Events) and the developments in the US Multiple input Tracking and Control System (MTRACS).

Operating platform:
MRTDF - VMS, DATAFUSE - UNIX/VMS.

Code language:
MRTDF - FORTRAN, DATAFUSE - C.

Input data - tracks:
MRTDF and DATAFUSE accept slant tracks as the required input. The data fields in slant coordinates group range, apparent azimuth, range rate (from Doppler), azimuth rate(MRTDF?). For both(?) approaches the SNR is not yet used in the fusion of tracks.

Input data - coordinate conversions:
In order to align data to a common coordinate system when testing hypotheses, coordinate transformations are employed by MRTDF and DATAFUSE. In MRTDF the appropriate transformations are read from the Coordinate Registration Tables (CRTs) which are updated approximately every 12(?) minutes for each Dwell Illumination Region (DIR). For the case of DATAFUSE being developed at 1RSU, the transformations will be provided by the Coordinate Registration System. Other ways of implementing DATAFUSE can be configured by using virtual ionospheric heights to define the propagation modes.

Input data - uncertainty in track data and coordinate conversions:
The management of uncertainty in the input data is clearly an area for improvement in both approaches. MRTDF accepts the reported uncertainty in the slant track but approximates the uncertainty in the ground track due to the coordinate conversion. DATAFUSE presently approximates the track and coordinate conversion uncertainties as a single uncertainty in slant coordinates.

Coordinate space for comparison of track data:
MRTDF performs the comparison of track data by transforming all slant tracks into ground coordinates. DATAFUSE uses a generic approach by transforming the data into sensor coordinates which, in the case of two OTHR measurements, is slant coordinates.
To illustrate the difference in the two approaches, let us consider the case where we wish to compare OTHR slant track A with OTHR slant track B. Furthermore, consider the case where we wish to decide whether A is propagating on ionospheric mode EE and B is propagating on ionospheric mode FF. The approach of MRTDF is to transform A into ground coordinates using the conversions appropriate for mode EE and to transform B into ground coordinates using mode FF. The two resulting ground tracks are then compared in ground coordinates. The approach of DATAFUSE is to transform A into ground coordinates using mode EE and to then take the ground track to predict a slant track using the coordinate conversions appropriate for mode FF. A comparison between the slant track prediction from A and the actual slant track B is then made. In practice the comparison is performed in both possible slant coordinates.

Temporal comparison:

MRTDF compares tracks which have temporal overlap in addition to tracks which are closely separated in time. DATAFUSE compares tracks which in general exist only simultaneously.

Time alignment:

In order to compare data with different time stamps, it must first be aligned in time. In MRTDF the slant tracker is used to extrapolate all the data with a Kalman filter to the integer minute. In DATAFUSE linear interpolation and extrapolation are performed when appropriate to align data to the time stamp of the measurements.

Measure of closeness:

MRTDF and DATAFUSE perform a pair-wise comparison of the data. In MRTDF a Chi-squared test is used to determine whether or not ground track data is associated. The four ground coordinate parameters which are considered are (latitude, longitude, course and speed). In DATAFUSE a Mahalanobis distance is calculated for each pair of slant tracks; one track being that which is predicted and the other being that which is actually observed. The slant coordinate parameters which are considered are (slant range, slant azimuth, slant range rate and slant azimuth rate).

Association test for pair-wise comparison:

In MRTDF the Chi-squared test is used to determine whether pairs of tracks are associated. In DATAFUSE a threshold level for the Mahalanobis distance can be set. A track-mode pair having a Mahalanobis distance which is less than the threshold level is deemed to be associated.

Grouping associated tracks:

In the MRTDF nomenclature a group of associated tracks is known as a family with the constituent slant tracks being members of the family. All of the propagation modes listed by the CRTs are used to form the track family for less than (or equal to?) three family members. After three members have been linked to a track family, the modes used to check for additional members is a reduced list (* requires clarification & expansion). In DATAFUSE a group of tracks can be formed based on the pair-wise comparisons.

Track history:

MRTDF performs the comparison of tracks at the minute interval only without considering previous comparisons. If two tracks converge momentarily so that their parameters have the required instantaneous values for association, then MRTDF will effect an association so that both tracks are members of the same family. If the two tracks then diverge, MRTDF will cause the family to split so that the tracks regain their original labels. In contrast to MRTDF, DATAFUSE stores a history of the comparison of tracks so that the correlation scores from previous times influence the decision whether tracks are associated at a given time instant. A factor can be applied so that recent comparisons are weighted strongly thereby reflecting the expected improvements in track and CR accuracies as time advances.

Combination of associated tracks:

Once a group of tracks have been deemed to be associated, the final stage in the fusion process is that of combining the data to provide a new, enhanced estimate. In MRTDF, the ground tracks which successfully meet the Chi-squared test criteria are combined in a weighted average to produce a fused ground track. The uncertainties in the constituent ground tracks are used to determine the corresponding weights. When more than one track contributes to a single ground position, a search is made for the minimum ground range variance over all members of the track family. This value and the corresponding ground azimuth variance for the same track are then used for the uncertainty in the fused track. (check)

In DATAFUSE, the intended operation is to take those ground tracks which are formed from associated slant tracks and to combine the ground tracks using their uncertainties. The fused ground track will be the result of a weighted average of the ground tracks and the fused track will have an uncertainty which is essentially the intersection of the individual uncertainties.

Output:

In both MRTDF and DATAFUSE the output consists of slant tracks which are deemed to be associated and the corresponding mode assignments. Fused ground tracks and their uncertainties are currently produced by MRTDF and it is intended that DATAFUSE will soon have this capability.

Efficiency considerations:

A comparison of all tracks reported by an OTHR with all other tracks is a computationally expensive operation. In MRTDF the set of tracks with which to compare a given track is reduced by setting a gate (check). The approach of DATAFUSE is to establish a partition grid in slant space and to compare a track in a particular grid cell only with those tracks lying in neighbouring cells.

4.5 SECTION REFERENCES

- Yssel, W. 1994. "OTH Sensor Fusion," NRaD ROTH White Paper, San Diego, CA, (September).
- Shaw, S., J. Arnold, G. Ravichandran, and E. Lyon III. 1998. "Approaches to Detection, Tracking, and Fusion for Multiple OTH Radars," SRI International, TR 6631, Menlo Park, CA, (February).
- Shellshear, A. and I. Dall. 1995. "Implementation Issues for Datafuse at JFAS." DSTO White Paper (draft), Salisbury, AS, (October).
- Shellshear, A. 1995. "Recent Improvements to Datafuse," DSTO White Paper (draft), Salisbury, AS, (October).
- Johnson, W. and D. Kewley. 1996. "ALICE: A Fuzzy Logic Approach to Automatic Multi-Sensor Data Fusion," *Proc. 1st Internat. Discourse on Fuzzy Logic and the Management of Complexity (FLAMOC)*, pp. 278-282, Sydney, AS, (January).
- White, K. 1996. "MRTDF and DATAFUSE Comparison for MOA," private communication, DSTO Salisbury, AS, (September).

5. STATISTICAL SUMMARY OF ROTH R DATA FUSION

5.1 MRTDF ACCEPTANCE TESTING

Fusing data from multiple OTH radar can mitigate the technical challenges mentioned in section 3. A fusion algorithm will be able to detect CR errors and compensate for them because it has access to target information from multiple sensors. Next, a fusion algorithm will improve the accuracy of tracked targets for the following two reasons. First, it has target data from more than one sensor so it can average out the noise more effectively. Second, if the sensors provide geometrically diverse measurements (e.g., if the range component of one sensor is orthogonal to the other sensor and the measurements have different in-range and cross-range measurement uncertainties), then the fusion algorithm will provide tracking uncertainties which will be the intersection of the uncertainties of the individual sensors.

To quantify the performance of the MRTDF algorithm, several measures of effectiveness (MOEs) have been used in our analysis. These will be discussed in section 5.2. These are probability of multi-track, course and speed accuracy, target position accuracy, containment, and track maintenance. Note that the MOEs can be associated with the technical challenges as follows: (a) Mode Identification—probability of multitrack, position accuracy, and containment; (b) Coordinate Registration—course and speed accuracy and target position accuracy; (c) Range Resolution—containment and track maintenance.

5.1.1 MRTDF RESULTS

In June 1995, the ROTH Texas came online and underwent Navy Acceptance testing as a single ROTH system. In July 1995, it was brought together with the Virginia ROTH, and the Multiple ROTH Track Data Fusion (MRTDF) capability was installed. The performance of the combined operation of the Virginia and Texas ROTH systems was evaluated for several months in 1995.

In an effort to validate the operation of the MRTDF system and quantify the improvements brought about by the fusing of the data from the two ROTH systems, a Navy P-3 aircraft was tasked to fly 10 missions out of Jamaica in the coverage area of the two ROTH systems. The aircraft was instrumented with a GPS unit that was interfaced to a laptop computer to log the position, course, and speed of the aircraft as it flew specially designed geometries in the Caribbean. The data were processed and recorded in real time by each of the two ROTH systems working independently as single radars. These recorded data were later played through the systems connected in the "fused configuration," and the respective outputs were compared. Figure 5.1.1 shows an example of the MRTDF performance in tracking the P-3 aircraft. Several significant improvements in tracking are exhibited by the MRTDF algorithm when compared with the single system ROTHs, largely due to the reasons discussed above.

The MRTDF algorithm was also used for targets tracked during two periods: the first commencing 31 May 1995 at around 1730Z (Phase I), and the second, 15 September 1995 (Phase 2), commencing around 1800Z. We will discuss MRTDF tracks from Phase 1, which were matched with truth tracks provided by the Combined En-Route Radar Approach (CERAP) system located in Panama. This radar system ASR-9 (Airport Surveillance Radar) was built and installed by Westinghouse USA, and is currently operated by the government of Panama. Figure 5.1.2 shows the MRTDF track in comparison to the CERAP (truth) track as well as the single system ROTH tracks from the Virginia and Texas ROTH systems. These results were typical in our system, and we found that for the CERAP




Figure 5.1.1. P-3 Track for 31 May 1995 2230-2400Z.




Figure 5.1.2. CERAP Track for 1 June 1995 0001-0019Z.

data, the MRTDF algorithm performed according to system specifications. More tests are planned using enhanced tracker algorithms in the near future.

One of the expected improvements was better course and speed accuracy. Even though there were improvements in course and speed estimates with the MRTDF system compared to a single ROTH system, in isolated cases with difficult geometries, course and speed accuracy, averaged over the entire database, improved only slightly. In cases where multiple ionospheric modes were available, the MRTDF selected different sets of modes to assign to the radar slant tracks. This resulted in a nearly 25% reduction in the errors related to target positioning.

5.2 MEASURES OF EFFECTIVENESS FOR MRTDF

5.2.1 Introduction and Background

To quantify the performance of the MRTDF algorithm, several MOEs have been used in analysis of the data. These are probability of multitrack, position accuracy, area of uncertainty (AOU) and containment, course and speed accuracy, and track maintenance. As previously stated, MOEs can be associated with the technical challenges as follows: (a) Mode Identification—probability of multitrack, position accuracy, and containment; (b) Coordinate Registration—course and speed accuracy and target position accuracy; (c) Range Resolution—containment and track maintenance.

5.2.2 Probability of Multitrack

Multitrack in this instance refers to a single radar target that simultaneously produces more than one ground track on the radar. In the ideal OTH radar fusion system, all slant tracks originating from a single source would be associated, correlated, and fused into one single ground track and displayed to the operator. But since these slant tracks can be the result of radar returns from different ionospheric modes, from coverage areas using different operating frequencies, or from a physically separated radar, problems can occur. Some reasons that have contributed to the occurrence of multitracks in the past include: (a) a new track being compared to a family (a collection of associated slant tracks) before its parameters have had a chance to settle out, (b) tracks being assigned to an incorrect ionospheric mode, and (c) the mode being used for propagation by the track was not thought to be active and hence was not included in the ionospheric model. An efficient data fusion algorithm should not increase the number of multitracks. For calculation of the probability of multitracks, one takes the time multitracks are present and divides this by the time that targets are being held in track.

5.2.3 Position Accuracy

A favorite MOE for position accuracy is miss distance. Miss distance is calculated as the great circle distance between the ROTH positional estimate and the true position of the target. Positional accuracy cannot be measured without an accurate source of ground truth data. For system evaluations, truth data for target aircraft are taken from on-board GPS receivers or are gathered from ground-based microwave radars situated in the vicinity of the targets. Since a sizable portion of the miss distance in OTH radars comes from range errors resulting from ionospheric uncertainties and miss identification of ionospheric modes, data from a second, independent, near orthogonal looking OTH should substantially reduce the miss distances on targets being held simultaneously by two OTH radars.

5.2.4 AOU and Containment

When the ROTH system provides a position estimate to the user, it also supplies an estimate of the accuracy of that position. This is done in the form of an area of uncertainty (AOU), an ellipse whose probability of expected containment is 86%. With the addition of a detection from an independent, near orthogonal, second radar system, one would expect the positional accuracy to increase, and as a consequence, the area of uncertainty to decrease. Since both the containment percentage and the AOU can change, it becomes a matter of normalizing the results in order to measure the amount of improvement achieved. For this reason, the concept of normalized miss distance is introduced. The miss distance is normalized by dividing it by the distance from the center of the ellipse to the edge of the ellipse moving along a great circle route in the direction of the truth location (see figure 5.2.1). In this representation, when the normalized miss distance is less than one, the truth position is contained in the uncertainty ellipse; when equal to one, the truth position lies on the ellipse boundary; and when greater than one, the truth position lies outside of the uncertainty ellipse. With the data in this form, one can uniformly expand or contract the set of AOU's until the containments of the two data sets are identical and make a comparison of the difference in the average or median areas of uncertainty. Similarly, the set of AOU's can be expanded or contracted until the average or median AOU's are identical in area allowing comparison of the percentage of targets contained within the borders of these uncertainty ellipses.




Figure 5.2.1. A is the center of the ROTH AOU. C is the true position of the target aircraft. Normalized Miss Distance is AC divided by AB.

5.2.5 Course and Speed Accuracy

Target course and speed are critically important radar data characteristics. Since information can be gained both from the nature of the errors (bias) and from the distribution of the size of the errors, it is prudent to take signed as well as unsigned averages or medians of the course and speed errors. For the case of OTH radar, the radial component of the velocity can be taken directly from the Doppler shift of the returned signal, while the tangential component of velocity must be calculated from successive measurements of target azimuth. This fact gives rise to a velocity error profile that increases as the aspect of the radar to the aircraft approaches a broadside (low Doppler) geometry. Since target course is derived from these two velocity components, course error suffers the same error profile characteristic. With the addition of a second, near orthogonal radar, two highly accurate Doppler measurements can be obtained and a much more accurate estimate of both course and speed should be possible for targets being tracked simultaneously by two OTH radars. These inherent characteristics of OTH radar make these MOEs very appropriate for a fused OTH radar system. For the

MRTDF evaluation, average and median course and speed errors were calculated using the ROTHr estimates and data from the GPS and microwave truth data sources.

5.2.6 Track Maintenance

In the ROTHr system there are two areas of concern that relate to track maintenance. First, the target should be held and updated in a consistent and continuous fashion. Thus, fractional track holding time (FTHT) is calculated by taking the time a track was held and dividing it by the opportunity time for holding. The opportunity time for holding begins when the target has been detected and is beginning to be tracked and lasts until the target ceases to be a track candidate, i.e., leaves the coverage area, makes a significant maneuver, or lands. Track holding is considered continuous as long as there are no periods of track loss that exceed 3 minutes. If the track goes more than 3 minutes without an update, the track holding is penalized for the full duration of the missed detections. The second area of concern in track maintenance is that a given target be assigned a single track number and that it be known by that number for the duration of its holding. Thus, fractional track association time (FTAT) is calculated by taking the longest track holding with a single track number and dividing it by the total track holding as calculated for the numerator of the previous expression. For a system consisting of two independent OTH radars, one would expect significant improvement in target maintenance. Total holding time should increase since a second radar, using a different part of the ionosphere and operating at a different frequency, would be expected to experience different fading characteristics and could provide holding when the target might be momentarily lost from the first radar. As a consequence, with fewer losses in track, segments should be longer and be less likely to require relinking to the previous tracks in order to maintain a consistent track number. Track maintenance is calculated for all track segments corresponding to target tracks that can be associated with a truth source.

5.2.7 Summary and Conclusions

The set of MOEs described above, though not an exhaustive set, did allow for a quantitative evaluation of the ROTHr MRTDF performance. A comparison of the outputs from the fusion algorithm with the performance of each individual radar system (without the benefit of MRTDF), allowed for quantification of the gains attainable in each of the technically challenging areas mentioned earlier. Many different approaches to this fusion problem are yet to be tried and evaluated. Data fusion is beneficial to the system, and with an adequate set of MOEs, these gains can be quantified.

5.3 SECTION REFERENCES

- Yssel, W., W. Torrez, and R. Lematta. 1996. "Multiple Relocatable Over-the-Horizon Radar (ROTHR) Track Data Fusion (MRTDF)," *Proc. 9th Nat. Symp. Sensor Fusion*, pp. 37-45, Monterey, CA, (September).
- Yssel, W., W. Torrez, and R. Lematta. 1996. "Measures of Effectiveness for Multiple ROTHr Track Data Fusion (MRTDF)," *Proc. 1st Australian Data Fusion Symposium*, pp. 106-109, Adelaide, AS, (November).

6. THEORY AND PRACTICE OF MULTISOURCE-MULTISENSOR DATA FUSION

6.1 RECENT DEVELOPMENTS IN FUSING MICROWAVE RADAR TRACKS WITH RELO- CATABLE OVER-THE-HORIZON RADAR (ROTHR) TRACKS

Recent Developments in Fusing Microwave Radar Tracks with Relocatable Over-the-Horizon Radar (ROTHR) Tracks

William J. Yssel and William C. Torrez
Space and Naval Warfare Systems Center
San Diego, CA 92152-5001

Abstract - The objective of this paper is to discuss the application of the Advanced Tactical Workstation for the fusion of tracking data collected from independent surveillance systems. Two independent systems of particular interest for this discussion are the Relocatable Over-The-Horizon Radar (ROTHR) system operated by the US Navy, and the Multiple Input Tracking and Control Subsystem (MTRACS) which provides intelligence, surveillance and reconnaissance operations support for the United States Southern Command. The current MTRACS display provides the operator with a surveillance picture based solely on the kinematic aspects of the target measured by a net of microwave radars. In a future implementation, ROTHR tracks formed in regions of interest (but possibly outside the range of the microwave radars) will be fused with the microwave radar tracks. The work in this paper focuses on the application of the Advanced Tactical Workstation for the fusion of these disparate sensor tracks.

Keywords: over-the-horizon radar, target tracking, advanced tactical workstation

1 Introduction

Two independent systems of particular interest for this research are the Multiple Input Tracking and Control Subsystem (MTRACS) (cf. reference [1]) and the Relocatable Over-the-Horizon Radar (ROTHR) system (cf. references [2] and [3]). The current MTRACS display provides the operator with a surveillance picture based solely on the fused kinematic aspects of the targets measured by a net of microwave radars. In a future implementation, ROTHR tracks formed in regions of interest (but possibly outside the coverage area of the microwave radars) will be fused with the microwave radar tracks. This paper discusses the application of a commercial off-the-shelf (COTS) data processing and display system known as the Advanced Tactical Workstation (ATW) (based on the Over-the-Horizon Targeting Airborne Surveillance Information System (OASIS) offered by Raytheon-Texas Instruments) for the fusion of ROTHR tracks with microwave tracks produced by the MTRACS processor. As the primary user of the MTRACS and ROTHR systems, US SOUTHCOM is responsible for a vast surveillance region of operations. It is therefore para-

mount that products of these disparate sensor assets be fused for time-critical detection and monitoring activities. Operational requirements to be promulgated by the user in the near future may necessitate such fusion capabilities.

1.1 Relocatable Over-The-Horizon Radar

ROTHR is a long range, bi-static high frequency radar surveillance system tasked with tracking air targets over large land and ocean areas. The long ranges, together with the relatively low altitudes of the targets, require that the radar look beyond the line-of-sight, in other words, over-the-horizon (OTH). This is accomplished by refracting the signal through the ionosphere to points beyond the horizon. The radar transmitter is designed to radiate 200 kilowatts of RF power from 16 log-periodic, amplitude tapered array antennas. The ROTHR system surveys a coverage area consisting of azimuths and ranges up to 80 degrees and 2000 nautical miles. The US Navy operates two ROTHR systems (AN/TPS-71 ROTHRs) located in the United States. The first one, located in Virginia, looks south over the Gulf of Mexico and the Caribbean. The second one, installed in south Texas, looks to the

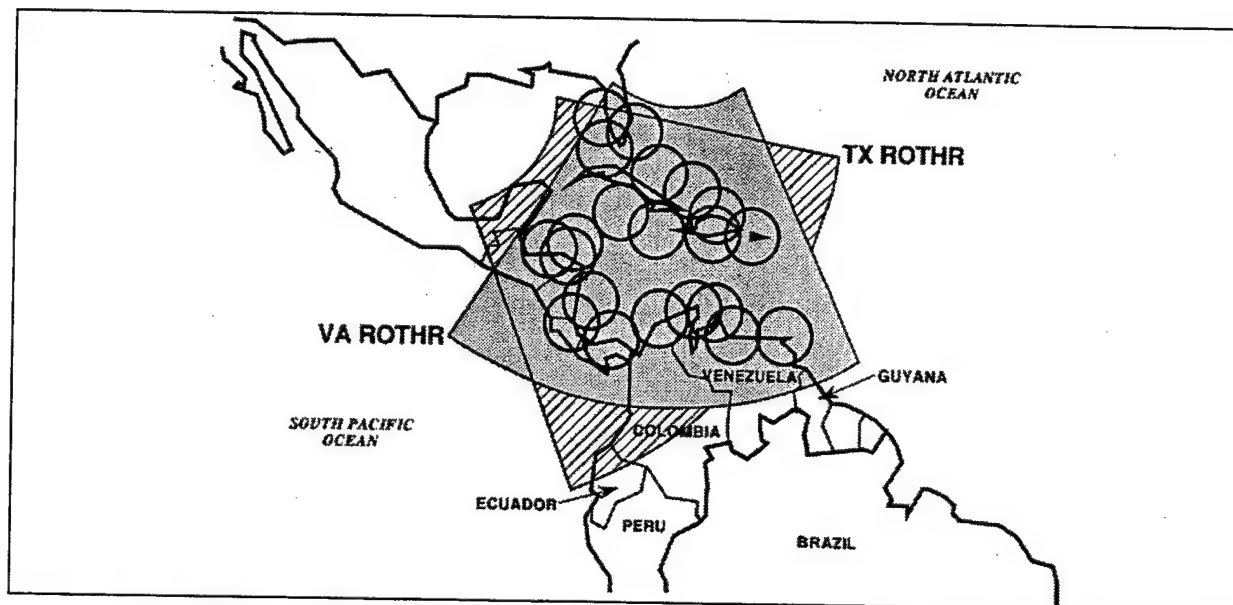


Figure 1. Geometry of ROTH and microwave radar coverage.

southeast and provides coverage of the Gulf and Caribbean. A notional geometry of the two ROTHs and the region of overlapping coverage with the microwave radars are shown in figure 1.

The primary mission of the ROTH system is to provide wide area surveillance in support of the Department of Defense's detection and monitoring activities in the Caribbean and Eastern and Pacific Regions. The ROTH system provides three features that complement the capabilities of the existing surveillance assets for detecting and tracking aircraft operating in the Gulf and the Caribbean. First, OTH radars have extremely large coverage areas (figure 1); second, with a suitable choice of additional OTH radar locations, they can be made to look at airspace deep into South America. Third, in contrast to the look-up capability of conventional ground radar, OTH radar provides a look down capability that provides better detection opportunities.

1.2 Caribbean Regional Operations Center

The Space and Naval Warfare Systems Center (SPAWARSCEN), San Diego, is designated as the lead field activity for upgrading the main-frame computers, communications, radars, and facilities for the Caribbean Regional Operations

Center, better known as CARIBROC. CARIBROC is located in Key West, Florida at Naval Air Station Key West. CARIBROC is a joint command, composed of representatives from the Navy, Air Force, Army, and Marine Corps.

The recently completed CARIBROC upgrade effort uses state-of-the-art technology in all development areas. The MTRACS system, a shore-based processing and display system, upgrades CARIBROC's current system. The updated system includes multi-sensor track correlation and multisource data fusion for up to 30 radars displayed on up to 32 consoles. It also interfaces with ROTH, tactical data links (4A, 11, and 16), and flight plan data, among many other enhanced features. The CARIBROC facility receives surveillance data from several sources including numerous ground based radar sensors, ROTH, as described previously, tethered aerostats, tactical data links, and special sources from within the United States Southern Command (US SOUTHCOM) areas of responsibility. The microwave data is processed, correlated, and fused for analysis at operator situation display consoles, but the ROTH data is kept separate and only displayed to the operators. Action is initiated commensurate with their analysis and mission instructions. While the primary mission of CARIBROC is to provide intelligence, surveillance, and reconnaissance (ISR) operation

support for US SOUTHCOM, their secondary mission is to provide counterdrug (CD) detection and monitoring support for the Joint Inter-Agency Task Force-East (JIATF-E).

2 Advanced Tactical Workstation

The Advanced Tactical Workstation, which has been hosted on a Pentium based personal computer, incorporates the Maneuvering Target Statistical Tracker (MTST) algorithm (cf. references [4] and [5]) to perform positional report filtering in support of track correlation processing, track solution generation, and target motion analysis. The MTST algorithm implements a statistical approach to perform Kalman filtering, backward smoother processing, and aberrant contact report identification to generate a set of track parameters, including primary solutions, based on areas of uncertainties (AOUs) that define a track based on the positional contact report data associated with the track. These primary solutions provide a statistical basis to support generation of geo-feasibility scores measuring the likelihood of association between existing tracks and individual positional reports; and the generation of estimated positions, AOUs, course, and speed for a target at any specified time. For the track data used in the demonstration described in section 3, the algorithm was provided with estimates of a target's location that contain velocity information in addition to the elliptical positional information. The input reports contain target course and speed, and the time of detection, the latitude and longitude of the center of the ellipse, and the semi-major and semi-minor axis lengths and the orientation of the ellipse.

2.1 Maneuvering Target Statistical Tracker

The Maneuvering Target Statistical Tracker (MTST) uses an iterated, extended Kalman filter algorithm to recursively incorporate a set of contact reports into updated track parameter estimates. A smoothing algorithm is used to support computation of track parameter estimates for any given time. A motion model that explicitly accounts for target maneuvering is used to predict target kinematics in the absence of observa-

tions. The update process will sort track events into time ordered sequence, will detect and remove aberrant reports as required from the updated solution and perform smoothing. The MTST forward filter processing incorporates new contact report data into the track solution estimates. MTST first takes the prior solution (the forward filter solution from the most recent event earlier in time than the new contact report) and if the motion model is being used, projects it forward to the time of the new report. This new solution is referred to as the prior-projected solution. Both the prior solution and the prior-projected solutions can be displayed simultaneously. A Kalman filter update is then done to combine the information from the prior projected solution and the new contact report. A forward filter solution, consisting of a state vector and covariance matrix, is computed from the updated information data. The forward filter solution at a given time represents the estimate of the target's location using only the data with report times up to and including that time. When a new report comes in, the forward filter processing is repeated for each new event. In the end, all of the events have updated forward filter solutions.

3 Demonstration and Discussion

Figure 2 below is a display of the 19 data contacts used in the fusion example. Barely visible are the AOU ellipses for several simulated microwave radar contacts (9 contacts) superimposed with both ROTH Virginia radar contact ellipses (8 contacts) and ROTH Texas radar contact ellipses (notice the difference in orientation of the 2 ellipses based on the ROTH Texas radar boresight, cf. figure 1). In figure 2, taken from an ATW display, the Track Display window is shown, and allows the user to view the data in multiple formats. In this display, we have chosen to view the AOUs of All Events. In subsequent figures, we will show the MTST Solution for the fused or merged track history.

Figure 3 is a close up of the southernmost microwave radar contacts (ending at 1920Z hrs), which were not clearly visible in figure 2. Notice the disparate size of the error ellipses for the simulated microwave ellipses when compared with the ROTH radar contact ellipses.

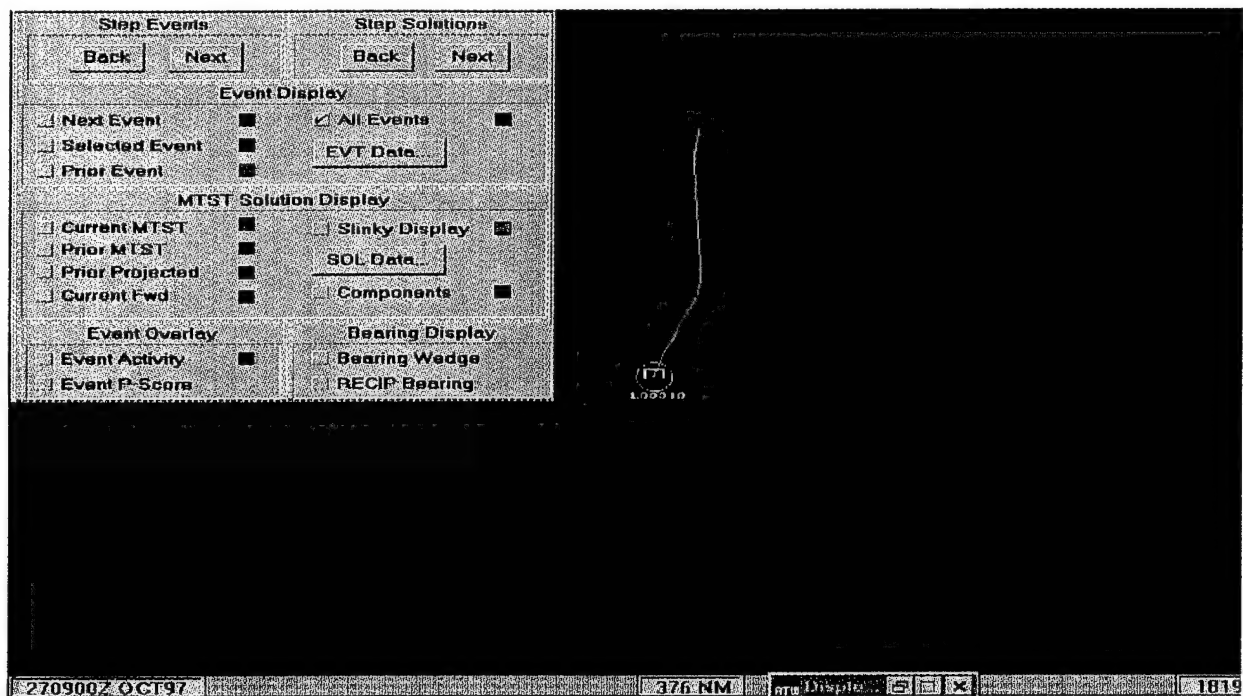


Figure 2. Area of uncertainty of 19 reports.

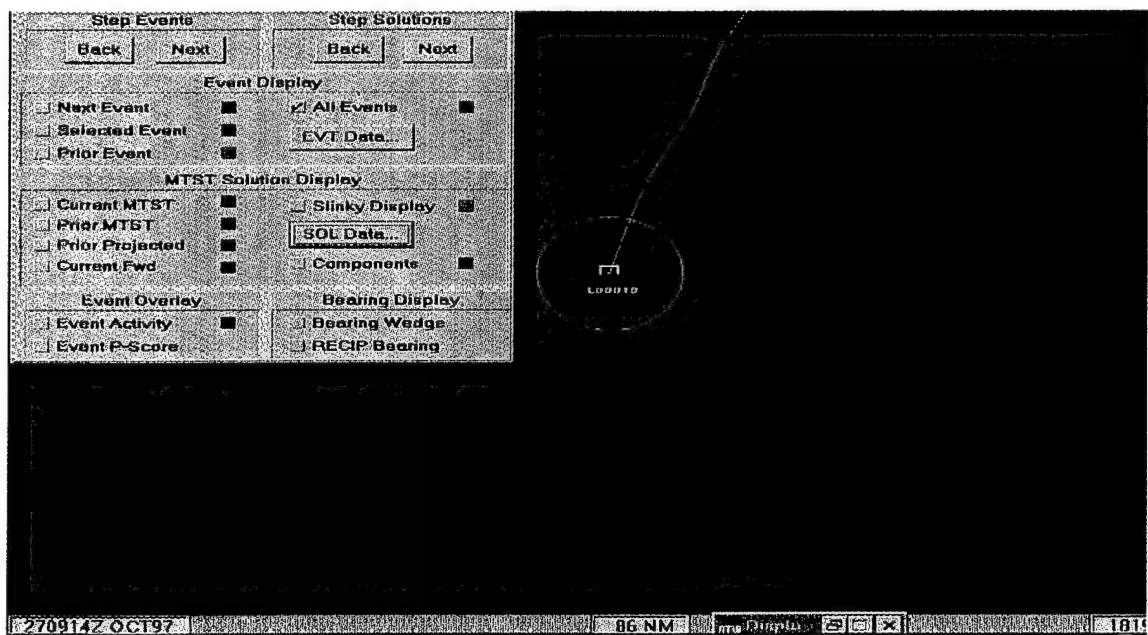


Figure 3. Microwave and ROTH reports near end of exercise.

the Prior Projected MTST solution (AOU ellipse denoted by line symbol xxx).

The image displays three screenshots from a military tracking system interface.

Top Left: M1ST Display

This window contains several menu sections:

- Step Event:** Back, New
- Step Solution:** Back, New
- Event Display:**
 - ☒ Next Event
 - ☒ Selected Event
 - ☒ Prior Event
 - ☒ All Events
 - ☒ Event Data
- M1ST Solution Display:**
 - ☒ Current M1ST
 - ☒ Prior M1ST
 - ☒ Prior Projection
 - ☒ Current CW
 - ☒ Slinky Display
 - ☒ SOL Data
 - ☒ Components
- Event Overlay:**
 - ☒ Event Activity
 - ☒ Event Store
- Bearing Display:**
 - ☒ Bearing Warden
 - ☒ REGIP Bearing

Top Right: Map

A map of the United States with a cluster of 'x' markers in the Pacific Ocean and a line of '+' markers.

Bottom: M1ST Solution Display

This window displays tracking data for a target on 18 SEP 97:

- DTG:** 18184600Z SEP 97
- POS:** 28N 079.05W
- CAS:** 1872.1 KTS
- MAJ:** 0.7 NM
- ANG:** 272.1
- MANUAL:** On
- 0.963 KTS**

Buttons for OK and Cancel are visible at the bottom.

6-6




Figure 6. Virginia ROTH hit at 1846Z hrs.

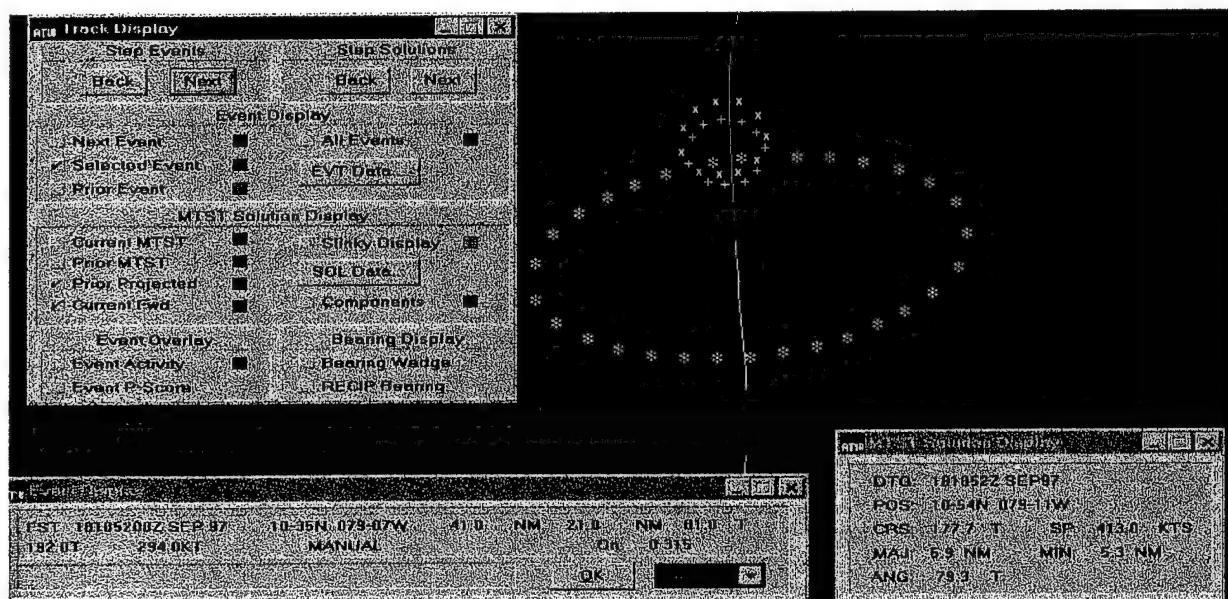


Figure 7. Virginia ROTH hit at 1852Z hrs.

Note that the Current Fwd solution (with AOU ellipse denoted by line symbol +++) is approximately the same size as the AOU ellipse for the given contact (Selected Event).

In figure 6, the selected event at the time 1846Z hrs, a ROTH VA contact, coincides in time with a microwave contact. The Current Fwd MTST filter solution (the AOU ellipse denoted by line symbol +++ , but covered slightly by the Prior Projected AOU ellipse denoted by line

symbol xxx) is still an ellipse which basically is that of the microwave contact at 1846Z hrs (see previous figure). Thus the MTST filter gives the solution based on the more accurate microwave contact, and behaves as we would expect a fusion filter to behave when fusing contacts from such disparate sensors.

In figure 7, the Current Forward solution ellipse (the AOU ellipse denoted by line symbol +++) has increased in size, but certainly not to

the size of the selected event at this time (a ROTHr VA contact at time 1852Z hrs whose ellipse is denoted by ***). This again shows that the MTST filter is fusing disparate products, but assigning a result based on the more accurate microwave contact reports.

Figure 7 displays the major benefit of fusing the microwave radar track with the ROTHr track for, in this case, the prior history from the microwave radar allows the current forwarded position to refine the selected ROTHr event posit to the north portion of the ellipse and greatly reduce the AOU that would otherwise be associated with the ROTHr only track, while maintaining the probability of containment.

Figure 8 shows that, at the time of the Selected Event, 1913Z hrs, the MTST filter fuses the microwave contact with the previous ROTHr contacts, and returns the Current Fwd ellipse, denoted by +++, which is basically, again, the smaller microwave ellipse. Note that the Prior Projected ellipse is largely commensurate with the previous fused ellipse.

4 Conclusions

In this paper, we have considered the MTST tracker, as implemented in the Advanced Tactical

Workstation for data that has both ROTHr and microwave radar contact data. It has been shown that the MTST solution gives essentially the microwave position and ellipse as the output (i.e. merges the ROTHr and microwave tracks overlapping in time) and that the position and ellipses of the merged tracks are essentially those of the microwave's. Furthermore, it has been demonstrated that if we move onto the portion of data which has ROTHr only inputs, the promulgated position of the target based on the past history derived from the microwave position is the union of the promulgated state with the large ROTHr position ellipse. Consequently, with fusion, the target state generated by the accurate microwave data serves to refine the positional estimate of ROTHr's large AOU, and provides a more accurate position and a significantly smaller AOU without decreasing the probability of containment. When MTST regains the target on microwave (probably a track jump), the microwave position and area of uncertainty (AOU) again dominate the state. Thus, it can be concluded that the ATW MTST fusion process does not sacrifice the containment of the track since the average AOU size is much reduced over the ROTHr only set of ellipses.

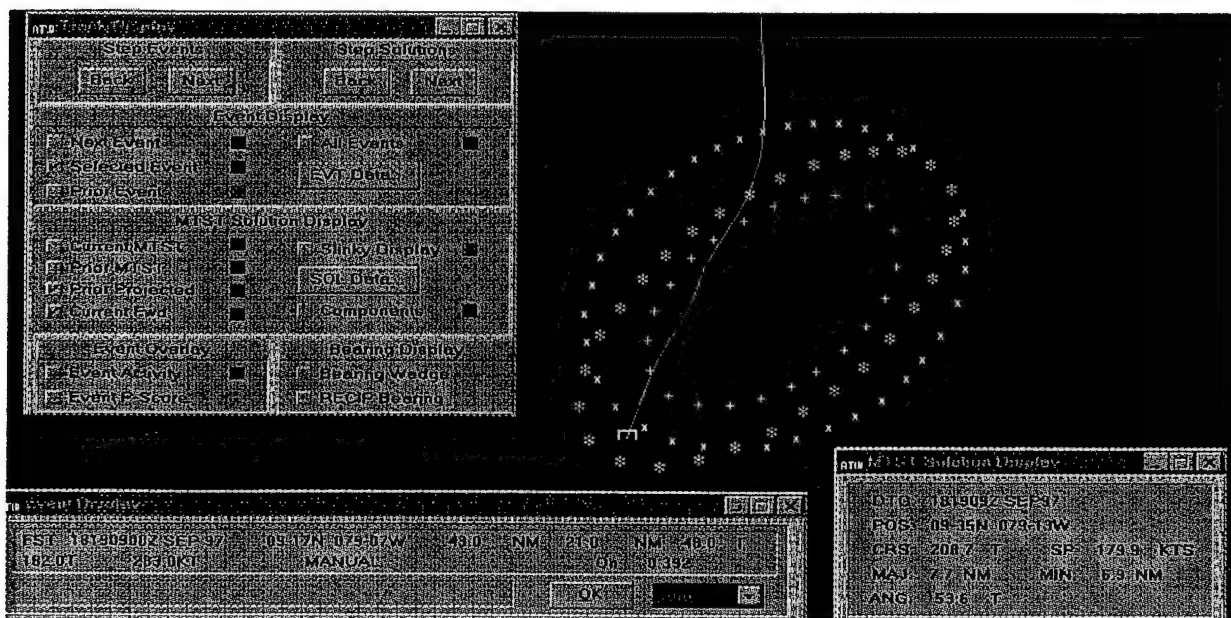


Figure 8. Selected Event at 1913Z hrs with Current Forward ellipse and Prior Projected ellipse.

Acknowledgements

The authors would like to thank the ROTHF Program Office, Naval Space Command, Dahlgren, VA, for their support of this work. They would also like to thank Mr. Robert Musil, ROTHF Project Leader at Space and Naval Warfare Systems Center, San Diego, CA for his support and encouragement.

References

- [1] W. Torrez and W. Yssel. Performance characteristics of data fusion methods with application to independent surveillance systems. *Proc. National Symposium on Sensor and Data Fusion*, 2:313-320, Oct 1997.
- [2] W. Yssel, W. Torrez, and R. Lematta. Measures of effectiveness for multiple ROTHF track data fusion (MRTDF). *Proc. 1st Australian Data Fusion Symposium*, 106-109, Nov 1996.
- [3] W. Yssel, W. Torrez, and R. Lematta. Multiple relocatable over-the-horizon radar (ROTHF) track data fusion (MRTDF). *Proc. 9th National Symposium on Sensor Fusion*, 37-45, Oct 1996.
- [4] C. Y. Chong, S. Mori, and K. C. Chang. Distributed multitarget multisensor tracking. In *Multitarget Multisensor Tracking: Advanced Applications*. Y. Bar-Shalom, Ed., Artech House, Norwood, MA, 247-295, 1990.
- [5] S. Mori, K. A. Demetri, W. Barker, and R. Lineback. A theoretical foundation of data fusion – generic track association metric. *Proc. 7th Joint Service Data Fusion Symposium*, 585-594, 1994.

6.2 TRACK FUSION AND COORDINATE REGISTRATION FOR MULTIPATH OVER-THE-HORIZON RADAR

The objective of the DSTO sensor fusion work is to determine the number of targets within the surveillance coverage and to estimate the state of each target. A multiple hypothesis approach has been used which accounts for ambiguity and uncertainty in the data. This approach, in which multiple hypotheses are recursively constructed and evaluated, provides a rigorous framework for examining less optimal techniques that may be suitable for operational implementation.

A multipath track fusion algorithm has been developed in which OTHR tracks and Coordinate Registration (CR) advice are used to construct feasible track-to-target hypotheses. The algorithm considers the uncertainties present in the radar track estimates and the CR transformations to ground coordinates. The CR uncertainty consists of a discrete component corresponding to the uncertainty in the utilization of a propagation path and a continuous component which describes the uncertainty in the ground state given that a propagation path is utilized.

The multipath track fusion algorithm has been described in three recent papers which have been included here in their entirety. The algorithm was first presented at the July 1997 SPIE conference in San Diego (*cf.* page 6-9). Path-dependent hypotheses are recursively constructed by considering each radar track state in turn to be associated with a propagation path and a given target. For a given hypothesis, target states represented by probability density functions are computed and a probability is evaluated for each hypothesis. A subsequent paper presented at the 1997 DSTO/AFOSR conference in Victor Harbor, South Australia, (*cf.* page 6-21), described an efficient strategy for pruning path-dependent hypotheses a priori given the marginal ordering of radar tracks and propagation paths. Finally, the implementation of recursive hypothesis evaluation and target state estimation was described in the paper presented at the 1998 SPIE conference in Orlando (*cf.* page 6-30). The evaluation of hypotheses in the presence of non-deterministic target state evolution was examined and the dependence of radar tracks due to common process noise, common measurements and common propagation path segments was identified.

The multipath track fusion algorithm has been implemented and integrated into a track fusion testbed for research and evaluation using simulated and real data. The track fusion testbed can also be used to accept data in real-time from the operational system at 1RSU.

Future work will include the extension of the multipath track fusion algorithm to incorporate data from microwave radars. This approach should enhance the association of propagation paths to OTHR tracks by utilizing associated microwave radar tracks which can be considered to arise via a single propagation path only. In addition, the accurate microwave radar data may be used to improve the CR transformations in the vicinity of overlapping microwave and OTH radar coverage.

Multipath Track Fusion for Over-the-horizon Radar

D.J. Percival and K.A.B. White,
High Frequency Radar Division,
Defence Science and Technology Organisation,
P.O. Box 1500, Salisbury, S.A. 5108, AUSTRALIA.

ABSTRACT

Over-the-horizon skywave radar exploits ionospheric propagation of HF signals to detect targets beyond the line-of-sight horizon. Multiple propagation paths between the radar sites and the target are often encountered, giving multiple resolved detections for a single target. An algorithm for the fusion of multipath tracks is outlined here which accounts for uncertainty in the coordinate registration transformation to ground coordinates. A multihypothesis track association procedure is described which may be appended to existing radar coordinate tracking filters. The probability for each feasible track association hypothesis is computed, and fused estimates for target states in ground coordinates are evaluated for each hypothesis.

Keywords: multipath track fusion, multihypothesis tracking, over-the-horizon radar

1. INTRODUCTION

Over-the-horizon skywave radar (OTHR) exploits the ionospheric refraction of HF signals to provide wide area surveillance of air and surface targets beyond the line-of-sight horizon.^{1,2} The complex ionospheric propagation environment³ may give rise to several propagation paths between the target and the radar site, leading to multiple resolved detections for a single target. A schematic diagram illustrating multipath propagation is shown in Figure 1, where one-hop propagation to a single target via two simple ionospheric layers may generate four return signals. Tracking filters are implemented in OTHR systems to associate target detections in the presence of clutter and to estimate states for multiple targets in radar coordinates.^{4,5} If multipath propagation conditions exist, such a tracking filter will generate multiple tracks for a given target.

To resolve the multipath ambiguity and to estimate true target states in geographic coordinates, a knowledge of the utilised ionospheric propagation path for each radar track is required. Possible propagation paths may be determined by numerical ray tracing through an ionospheric model, where the free parameters in the model are estimated using prior information and real-time ionospheric sounding measurements.^{6,7} The transformation of radar coordinate track states to physical ground coordinates, after identifying the utilised propagation path, is known as *coordinate registration* (CR).⁸ Multipath tracks may be identified when the corresponding transformed tracks to ground coordinates superimpose to form a single track. However, sets of apparently correlated tracks which are multipath candidates may also be due to multiple targets traversing similar but displaced trajectories. Since the association of propagation paths to radar tracks is not known *a priori*, several possible path association hypotheses must be examined by testing for track association after applying plausible CR transformations to each candidate track.

The association and fusion of OTHR multipath tracks has been the subject of several research efforts, most of which have worked with the output of existing tracking filters implemented in radar coordinates. This suboptimal hierarchical approach yields algorithms which are compatible with established tracking architectures, and are computationally attractive.⁹ A multipath track association algorithm has been developed in support of Australia's Jindalee OTHR project which evaluates a track association metric in a pairwise comparison of OTHR tracks.^{10,11} Pattern recognition algorithms have been applied to determine associated OTHR tracks using patterns which characterise multipath tracks of the same target.^{12,13} More recently, a multipath probabilistic data association (PDA) filter has been developed for OTHR which tracks targets in ground coordinates directly using multipath detections, given deterministic CR advice for each propagation path in a simple ionospheric model.^{14,15} Multipath track fusion

Corresponding Author: John Percival; Email: john.percival@dsto.defence.gov.au




Figure 1. Multipath propagation between an OTH radar at point O and a target at point P for a simple ionosphere consisting of two spherically symmetric reflecting layers. Four propagation paths are possible, namely transmission and reflection via the F layer or the E layer (FF and EE paths respectively), or mixed layer propagation where the signal is transmitted via one layer and is received via the other layer (EF and FE propagation). While FF, EE and mixed (EF and FE) returns may be resolved using their distinct propagation path lengths, the two mixed paths (EF and FE) may be separated only on the basis of different apparent target azimuth.

algorithms^{16,17} have been extended to multiple radar systems. A maximum likelihood approach to multipath CR using hidden Markov modelling of the likelihood function is also reported in the literature.¹⁸

In this paper, an algorithm is outlined for the association and fusion of multipath tracks generated by existing OTHR tracking filters. All path-dependent track-to-target association hypotheses are recursively constructed, and the probability of each hypothesis and the fused ground track states within each hypothesis are evaluated. In Section 2, OTHR tracking and an ionospheric state estimation approach to CR are briefly reviewed, and a formalism is introduced which accounts for both propagation path uncertainty and CR transformation uncertainty in evaluating ground coordinate track states. In Section 3, a track-to-target hypothesis tree is introduced to formulate and enumerate all multipath track association hypotheses. The evaluation of hypothesis probabilities and fused track states is outlined, exploiting the similarity of the multipath track fusion problem with multihypothesis multisensor tracking reported in the literature.^{9,19}

2. OTHR TRACKING AND MULTIPATH COORDINATE REGISTRATION

In this section, the transformation of radar tracks to ground coordinates using uncertain coordinate registration (CR) advice is treated. In a multipath propagation environment, two types of CR uncertainty arise, namely a discrete uncertainty in which propagation path applies for a given radar track, and the continuous transformation uncertainty due to incomplete knowledge of the utilised propagation path. It is assumed that tracking in radar coordinates is performed using a Kalman filter based algorithm, where target states and covariances are estimated from associated detection measurements in a multitarget and cluttered environment.⁴ Issues such as track initiation and deletion, track confidence, target manoeuvre modelling, tracking across radar dwell regions, clutter modelling and the resolution of velocity ambiguities are assumed to be addressed in the tracking filter. Also assumed is the provision of CR advice from ionospheric sounding measurements and modelling such that for a given radar track state estimate, a finite list of possible propagation paths and path probabilities can be identified, and for each path, the corresponding transformation to ground coordinates and a covariance describing the uncertainty in that transformation is provided. These requirements are detailed in the sections which follow.

2.1. Tracking in Radar Coordinates

Let $x_j(k) \in \mathcal{R}$ be the state associated with the j th track τ_j at time k in the radar coordinate space \mathcal{R} . In a multitarget cluttered environment, a set of states $\{x_j(k)\}_{j=1}^N$ is estimated from radar detection measurements using an established tracking filter.⁴ In general, the conditional probability density function (pdf) $p(x_j(k)|Z_j^k)$ for each state is determined, where $Z_j^k = \{z_j(i)\}_{i=1}^k$ is the set of measurements to time k associated with τ_j by the tracking filter. It is understood here that the interpretation of $z_j(k)$ depends on the choice of tracking filter and its data association algorithm. For example, using a nearest neighbour measurement to track association strategy, $z_j(k)$ is a single detection associated with τ_j ; using a PDA filter, $z_j(k)$ is interpreted as a set of detections lying within the filter measurement gate. The specific choice of tracking filter is important when considering track dependencies in the subsequent multipath track fusion stage (Subsection 3.3).

In Kalman-type trackers, the conditional pdf $p(x_j(k)|Z_j^k)$ has a Gaussian distribution, and the sufficient statistics computed are the conditional mean $\hat{x}_j(k|k) = E\{x_j(k)|Z_j^k\}$ and covariance $P_{x_j}(k|k) = \text{cov}\{x_j(k)|Z_j^k\}$, so that

$$p(x_j(k)|Z_j^k) = \mathcal{N}(x_j(k); \hat{x}_j(k|k), P_{x_j}(k|k)), \quad (1)$$

with

$$\mathcal{N}(x; \hat{x}, P) = |2\pi P|^{-1/2} \exp \left[-(1/2)(x - \hat{x})' P^{-1} (x - \hat{x}) \right] \quad (2)$$

the frequently used notation for the multivariate Gaussian distribution, and where the prime in (2) denotes the matrix transpose. The output of the tracking filter is a set of tracks $\{\tau_j\}_{j=1}^N$ in radar coordinates, with the track τ_j associated with the sequences of state estimates $\{\hat{x}_j(k|k)\}_{k=1}^K$ and covariances $\{P_{x_j}(k|k)\}_{k=1}^K$, when the track is maintained for K sample times.

2.2. Ionospheric State Estimation

Coordinate registration of radar track states to ground locations requires the ionospheric propagation paths between the target and the radar sites to be specified. In a multipath propagation environment, multiple resolved detections of a single target give rise to several tracks in radar coordinates; the correct propagation path for each track must be identified before the required CR transformation is applied. Knowledge of the signal propagation path may be obtained by numerically tracing rays through a parametric model of the ionospheric electron density profile over the radar coverage region.

Ionospheric modelling from sounding data may be cast as a state estimation problem, where the ionosphere at time k is described by a state vector $w(k)$, the components of which are parameters in a dynamic model of the ionospheric electron density. The dimension of $w(k)$ is determined by the number of free parameters in the ionospheric model chosen to yield the required CR transformation accuracy. In general, an ionospheric state estimation procedure may be used to compute the conditional pdf $p(w(k)|D^k)$, with $D^k = \{d(0), \{d(i)\}_{i=1}^k\}$ collectively denoting the prior physical information $d(0)$ and the sequence of ionospheric sounding measurements $\{d(i)\}_{i=1}^k$. Kalman-based estimation procedures yield the conditional mean $\hat{w}(k|k) = E\{w(k)|D^k\}$ and covariance $P_w(k|k) = \text{cov}\{w(k)|D^k\}$ from the sounding measurements and physical constraints. In Ref. 6, an extended Kalman filter is employed to estimate parameters in a simple single-layer ionospheric model from noisy sounding data asynchronously recorded by a geographically distributed array of vertical sounders. In Ref. 7, an H_∞ -type filter is applied to the same problem, to explore robust estimation procedures when model uncertainties and malicious noise are present. Ionospheric models with multiple layers are required to describe the multipath propagation considered here. In the sequel, it is assumed that an ionospheric state estimation procedure is available so that real time CR transformation and propagation path advice is provided.

2.3. Multipath Coordinate Registration

Numerical ray tracing through the ionosphere described by the state $w(k)$ is used to compute the CR transformation $T^{(m)} : \mathcal{R} \rightarrow \mathcal{G}$ from radar coordinates \mathcal{R} to ground coordinates \mathcal{G} for each feasible propagation path $m \in \mathcal{M}$, where \mathcal{M} is a set of integers which label the feasible propagation paths. For the state $x_j(k) \in \mathcal{R}$ associated with the track τ_j , the corresponding plausible multipath states in ground coordinates are given by

$$y_j^{(m)}(k) = T^{(m)}(x_j(k), w(k)) \in \mathcal{G}, \quad m \in \mathcal{M}(x_j(k), w(k)). \quad (3)$$

For the sake of simplicity, the total number M of feasible propagation paths at all times k is assumed to be known *a priori*, so that $\mathcal{M}(x_j(k), w(k)) \subseteq \{1, 2, \dots, M\}$, for all $x_j(k)$ and $w(k)$. In practice, the radar coordinate space \mathcal{R} is quantised so that $\mathcal{R} = \cup_i \mathcal{R}_i$ with quantisation cells \mathcal{R}_i . For a given ionospheric state $w(k)$, the CR transformations $T^{(m)}(x_j(k), w(k))$ are then computed for each quantisation cell \mathcal{R}_i and for each propagation path $m \in \mathcal{M}(x_j(k), w(k))$, with $x_j(k) \in \mathcal{R}_i$. The resulting transformations for each \mathcal{R}_i are tabulated in a *coordinate registration table*.

In general, the CR transformation $T^{(m)}$ is a nonlinear function of its arguments. However, for many propagation conditions and for most target states, $T^{(m)}$ may be linearised about the current state estimates. Expanding (3) to first order then gives

$$y_j^{(m)}(k) = \hat{y}_j^{(m)}(k|k) + \delta y_j^{(m)}(k), \quad (4)$$

with

$$\hat{y}_j^{(m)}(k|k) = T^{(m)}(\hat{x}_j(k|k), \hat{w}(k|k)), \quad m \in \mathcal{M}(\hat{x}_j(k|k), \hat{w}(k|k)), \quad (5)$$

and where

$$\delta y_j^{(m)}(k) = J_{x_j}^{(m)}(k) [x_j(k) - \hat{x}_j(k|k)] + J_w^{(m)}(k) [w(k) - \hat{w}(k|k)]. \quad (6)$$

In writing (6), the transformation Jacobians $J_{x_j}^{(m)}$ and $J_w^{(m)}$ are respectively given by

$$J_{x_j}^{(m)}(k) = \left. \frac{\partial T^{(m)}(x_j, w)}{\partial x_j} \right|_{(\hat{x}_j(k|k), \hat{w}(k|k))}, \quad J_w^{(m)}(k) = \left. \frac{\partial T^{(m)}(x_j, w)}{\partial w} \right|_{(\hat{x}_j(k|k), \hat{w}(k|k))}. \quad (7)$$

Substituting (6) into the ground state covariance

$$P_{y_j}^{(m)}(k|k) = E\{\delta y_j^{(m)}(k) \delta y_j^{(m)}(k)' | Z_j^k, D^k\} \quad (8)$$

gives

$$P_{y_j}^{(m)}(k|k) = J_{x_j}^{(m)} P_{x_j}(k|k) J_{x_j}^{(m)'} + P_T^{(m)}(k|k), \quad (9)$$

with

$$P_T^{(m)}(k|k) = J_w^{(m)} P_w(k|k) J_w^{(m)'} \quad (10)$$

defined here as the transformation covariance. In deriving (9), the independence of the $x_j(k)$ and $w(k)$ is assumed.

To recap, for a given state estimate $\hat{x}_j(k|k) \in \mathcal{R}$ corresponding with the radar coordinate track τ_j , and for a given estimate $\hat{w}(k|k)$ of the ionospheric state, a ground coordinate state estimate $\hat{y}_j^{(m)}(k|k) \in \mathcal{G}$ may be derived using (5) for each possible propagation path $m \in \mathcal{M}(\hat{x}_j(k|k), \hat{w}(k|k))$. The uncertainty in this estimate for the linearised CR transformation is described by the covariance $P_{y_j}^{(m)}(k|k)$ given by (9). There are two contributions to the ground state covariance $P_{y_j}^{(m)}(k|k)$, given by the two terms on the right side of (9). The first term gives the contribution of the radar state uncertainty described by the covariance $P_{x_j}(k|k)$ computed by the tracking filter, transformed to ground coordinates using the appropriate Jacobian $J_{x_j}^{(m)}$. In practice, $J_{x_j}^{(m)}$ may be approximated by the finite difference ratio

$$J_{x_j}^{(m)} \approx \frac{T^{(m)}(\hat{x}^{(i+1)}, \hat{w}) - T^{(m)}(\hat{x}^{(i)}, \hat{w})}{\hat{x}^{(i+1)} - \hat{x}^{(i)}}, \quad (11)$$

with $\hat{x}^{(i)} \in \mathcal{R}_i$, $\hat{x}^{(i+1)} \in \mathcal{R}_{i+1}$ and $x_j \in \mathcal{R}_i \cup \mathcal{R}_{i+1}$ in adjacent cells \mathcal{R}_i and \mathcal{R}_{i+1} of the CR table, and where the time arguments have been omitted for clarity. The second term on the right side of (9) is the transformation covariance $P_T^{(m)}(k|k)$ which describes the CR transformation uncertainty for the m th propagation path. $P_T^{(m)}(k|k)$ may be evaluated using (10) from the ionospheric covariance $P_w(k|k)$ computed by the ionospheric state estimation procedure, and from the Jacobian $J_w^{(m)}$ inferred from numerical ray tracing through the estimated ionosphere.

2.4. Multipath Ground State Density

The state estimate (5) and covariance (9) describe the distribution of the transformed track state in ground coordinates, for a given propagation path m . It remains to give a probabilistic description of path utilisation. Let $\theta_j^{(m)}$ denote the event that the radar track τ_j is propagating via the m th path. Then the conditional probability

$$\beta_j^{(m)}(k) = \Pr\{\theta_j^{(m)} | \tilde{Z}_j^k, \tau_j\}, \quad m \in \mathcal{M}(x_j(k), w(k)) \quad (12)$$

may be defined, with $\tilde{Z}_j^k = \{Z_j^k, D^k\}$ for convenience and where the normalisation

$$\sum_{m \in \mathcal{M}} \beta_j^{(m)}(k) = 1 \quad (13)$$

is used. It is assumed in what follows that $\beta_j^{(m)}(k)$ has been computed as part of the CR advice using prior physical information and numerical ray tracing through ionospheric models, and the resulting path probabilities tabulated for each cell \mathcal{R}_i , with $x_j \in \mathcal{R}_i$ and $m \in \mathcal{M}$.

Since the propagation paths are mutually exclusive, it follows from total probability that the conditional pdf for the j th ground state $y_j(k) \in \mathcal{G}$ is given by

$$p(y_j(k) | \tilde{Z}_j^k, \tau_j) = \sum_{m \in \mathcal{M}} \beta_j^{(m)}(k) p(y_j(k) | \tilde{Z}_j^k, \tau_j^{(m)}), \quad (14)$$

where $p(y_j(k) | \tilde{Z}_j^k, \tau_j^{(m)})$ is the pdf of (3), and with $\tau_j^{(m)} = \tau_j \cap \theta_j^{(m)}$ denoting the j th ground track when propagation is via path m . It is convenient to extend the summations in (13) and (14) to all $m \in \{1, 2, \dots, M\}$ by defining $\beta_j^{(m)}(k) = 0$ when $m \notin \mathcal{M}$ and $y_j^{(m)}(k)$ is not defined. For the linearised CR transformation, substitution of the Gaussian path conditional pdf

$$p(y_j(k) | \tilde{Z}_j^k, \tau_j^{(m)}) = \mathcal{N}(y_j(k); \hat{y}_j^{(m)}(k|k), P_{y_j}^{(m)}(k|k)) \quad (15)$$

into (14) and using the summation extension gives the Gaussian mixture density

$$p(y_j(k) | \tilde{Z}_j^k, \tau_j) = \sum_{m=1}^M \beta_j^{(m)}(k) \mathcal{N}(y_j(k); \hat{y}_j^{(m)}(k|k), P_{y_j}^{(m)}(k|k)) \quad (16)$$

for the ground state. When no additional ground truth or track association information is available, (16) provides a complete description of the ground state pdf corresponding with track τ_j , incorporating CR transformation and propagation path uncertainty. It follows from (16) that the ground state conditional mean $\hat{y}_j(k|k)$ and covariance $P_{y_j}(k|k)$ are respectively given by

$$\hat{y}_j(k|k) = \sum_{m=1}^M \beta_j^{(m)}(k) \hat{y}_j^{(m)}(k|k) \quad (17)$$

and

$$P_{y_j}(k|k) = \sum_{m=1}^M \beta_j^{(m)}(k) \left[P_{y_j}^{(m)}(k|k) + \hat{y}_j^{(m)}(k|k) \hat{y}_j^{(m)}(k|k)' \right] - \hat{y}_j(k|k) \hat{y}_j(k|k)'. \quad (18)$$

3. MULTIHYPOTHESIS MULTIPATH TRACK FUSION

In this section, a multihypothesis track association and fusion algorithm is described in which the probabilities of all track-to-target association hypotheses for all feasible propagation paths are evaluated, and the fused track states estimated for each hypothesis. Tracks are deemed to be associated if their corresponding states are estimated using detections from the same target propagating via different propagation paths. This algorithm draws on previous work in multisensor multihypothesis tracking,^{9,19} where an analogy may be drawn between resolved multipath and multisensor track association and fusion. For notational brevity, arguments are omitted throughout this section when the meaning is clear.

3.1. Track Clustering

The initial classification of radar tracks into clusters of possibly associated tracks greatly reduces the computational load of hypothesis evaluation. Cluster membership is dynamic as the associated tracks evolve. Two radar coordinate tracks τ_1, τ_2 are common elements in a track cluster at time k if their corresponding transformed ground states $y_1^{(m_1)}(k)$ and $y_2^{(m_2)}(k)$ may be estimated from detections of the same target, with m_1 and m_2 respectively labelling compatible propagation paths which satisfy $m_1 \neq m_2$ and any other physical constraints which may be imposed (e.g., Subsection 3.2). Rather than test all feasible propagation paths separately for a given track pair, an approximate test condition for common cluster membership is

$$(\hat{y}_1 - \hat{y}_2)' (P_{y_1} + P_{y_2})^{-1} (\hat{y}_1 - \hat{y}_2) \leq \gamma, \quad (19)$$

for a chosen threshold γ , where the state estimate (17) and covariance (18) for the Gaussian mixture model is used. In writing (19), it is implicitly assumed that any between-track dependence can be neglected when applying this membership condition. Clusters may be iteratively constructed by considering each track in turn, and either adding the track to an existing cluster if the condition (19) is satisfied for at least one existing cluster element, or forming a new distinct cluster. Cluster formation and management strategies are well developed in the context of multihypothesis tracking,²⁰ and are not pursued further here.

3.2. Hypothesis Formation

Track-to-target association hypotheses in a multipath environment may be formulated and enumerated using a *hypothesis tree*. Each path through the hypothesis tree represents one association hypothesis of all cluster tracks with targets, where a mutually feasible set of propagation paths is available. The hypothesis tree is constructed by considering a track cluster $\{\tau_j\}_{j=1}^N$, and assigning a target label t to each track such that every target has at least one track assigned to it. The hypothesis tree when $N = 4$ for a static cluster is shown in Figure 2. The tree is constructed from left to right as the tracks in the cluster are considered in turn. The tree root is formed with the first track τ_1 assigned to target $t = 1$, with the root node so labelled. If there are M feasible propagation paths, then there are M possible ground tracks $\{\tau_1^{(m)}\}_{m=1}^M$, corresponding to the M path-dependent track-to-target association hypotheses. If no other track is associated with target $t = 1$, then the target ground state is given by the mixture distribution (16) with $j = 1$ and $m = m_1$. False tracks are neglected here for the sake of simplicity, but may be accommodated by assigning false tracks to a "false target" (e.g., $t = 0$), for which the false detections may be propagating via any path.

Two hypothesis branches form when the second track τ_2 is considered, namely $t = 1$ if τ_1 and τ_2 are associated multipath tracks of the same target, or $t = 2$ if τ_2 is due to a new target. In the latter case, the track states are independent, and τ_2 is treated in the same manner as τ_1 . If τ_1 and τ_2 are associated, then there exists a path pair (m_1, m_2) with $m_1 \neq m_2$ such that the $t = 1$ target state can be estimated by fusing the individual ground state pdfs $p(y_1|\tilde{Z}_1^k, \tau_1^{(m_1)})$ and $p(y_2|\tilde{Z}_2^k, \tau_2^{(m_2)})$. Fused state estimation is outlined in Subsection 3.3. The specification of m_1 leaves $(M - 1)$ remaining paths for m_2 , so that there are $M(M - 1)$ path-dependent hypotheses for two associated tracks.

In general, there are $M P_n = M!/(M - n)!$ path-dependent hypotheses when n tracks are associated with a given target. However, a substantial reduction in this number can be achieved by considering cluster tracks in a prescribed *marginal order* when forming the hypothesis tree, so that physically impossible path combinations are eliminated *a priori*. Marginal ordering of a radar track cluster $\{\tau_j\}_{j=1}^N$ follows by ordering a given component of their corresponding state estimates $\{\hat{x}_j\}_{j=1}^N$. Since the transformed states $\{\hat{y}_j^{(m)}\}_{m=1}^M$ may also be marginally ordered for each \hat{x}_j , at best a constraint of the form $m_n < m_{n-1} < \dots < m_1$ may be derived for the association of tracks $\tau_1, \tau_2, \dots, \tau_n$ in that order. The details of this hypothesis reduction strategy depend on the accuracy and form of the CR advice, and are omitted here for the sake of brevity. However, it is straightforward to show that the number of path-dependent hypotheses reduces to $M C_n = M P_n/n!$ when this additional constraint is imposed.

The number of nodes at a given depth in a static hypothesis tree gives the number of mutually exclusive track-to-target association hypotheses, when propagation path considerations are neglected. At a given tree depth $j = J$, the




Figure 2. Track-to-target hypothesis tree for a cluster of four tracks, constructed from left to right as the tracks labelled $j = 1, 2, 3, 4$ are considered in turn. Each tree node is labelled with the target label $t = 1, 2, 3$ or 4 .

number of hypotheses assigning T targets with J tracks is given by the *Stirling number of the second kind* $S(J, T)$. The total number of association hypotheses is then

$$H(J) = \sum_{T=1}^J S(J, T), \quad (20)$$

which in the literature is called the *Bell number* or *exponential number*.²¹ The cases $H(1) = 1$, $H(2) = 2$, $H(3) = 5$ and $H(4) = 15$ are illustrated in Figure 2. The evaluation of $S(J, T)$ and $H(J)$ is the subject of Appendix A. When the set of feasible propagation paths is also considered, each node in the hypothesis tree then represents a set of path-dependent association hypotheses, which must be evaluated separately. While this may require more than $H(J)$ hypotheses be considered at tree depth J , it is also possible that the bound M on the number of propagation paths will reduce the number of feasible hypotheses when tree branches which imply $T > M$ are pruned.

The hypothesis tree construction may be generalised to include target dynamics, as illustrated in Figure 3. A simple cluster of three tracks in radar coordinates is shown in Figure 3(a), with tracks τ_1 and τ_2 represented respectively by state estimates $\hat{x}_1(k|k)$ and $\hat{x}_2(k|k)$ with $k = 1, 2$, while the track τ_3 is not initiated until $k = 2$, so that only $\hat{x}_3(2|2)$ is shown. The corresponding hypothesis tree is shown in Figure 3(b), where the tree depth is now indexed by the ordered pair (j, k) . In contrast with detection-to-track hypotheses in multihypothesis tracking,²⁰ the number of track-to-target hypotheses here remains manageable because estimates at different times for the same radar track must relate to the same target. That is, new estimates for a previously established track constrain the track-to-target hypotheses to those already under consideration. For modest track clusters (e.g., $N \leq 6$ so that $T \leq 6$), exhaustive evaluation of all association hypotheses is computationally possible.

3.3. Hypothesis Evaluation and Target State Fusion

A recursive algorithm may be derived to compute the probability of each path-dependent track-to-target association hypothesis, given the probability of the parent hypothesis on the preceding node along the same branch of the hypothesis tree. At the tree root node when the first track τ_1 is considered, there are M path-dependent association




Figure 3. (a) A cluster of three tracks τ_j , $j = 1, 2, 3$ in radar coordinates showing state estimates (filled squares) at times, $k = 1$ and $k = 2$. (b) The corresponding track-to-target hypothesis tree with hypothesis nodes (filled circles) labelled with a target identity number $t = 1, 2$ or 3 . Note that nodes corresponding with new estimates of established tracks have the same target label as all preceding nodes for that track.

hypotheses when M propagation paths exist, with each hypothesis corresponding to a ground track $\tau_1^{(m_1)}$, where $m_1 \in \{1, 2, \dots, M\}$. If λ_0 denotes one of these root hypotheses, then from Bayes theorem,

$$\Pr\{\lambda_0 | \tilde{Z}_1^k\} = \Pr\{\tau_1 \cap \theta_1^{(m_1)} | \tilde{Z}_1^k\} = C^{-1} \beta_1^{(m_1)}(k) L(\tau_1 | \tilde{Z}_1^k), \quad (21)$$

with the normalisation constant C incorporating the priors $\Pr\{\tau_1\}$ and $\Pr\{\tilde{Z}_1^k\}$, and where $L(\tau_1 | \tilde{Z}_1^k) \propto \Pr\{\tilde{Z}_1^k | \tau_1\}$ denotes the likelihood of the track τ_1 . Normalising (21) over all root hypotheses using (13) gives

$$\Pr\{\lambda_0 | \tilde{Z}_1^k\} = \Pr\{\tau_1^{(m_1)} | \tilde{Z}_1^k\} = \beta_1^{(m_1)}(k), \quad (22)$$

so that the root hypothesis probability is simply the probability that τ_1 propagates via mode m_1 .

The evaluation of an association hypothesis on an arbitrary node of the hypothesis tree proceeds by analogy with the treatment of hierarchical multihypothesis multisensor track fusion reported in the literature.^{9,19} Let $\bar{\lambda}$ denote the parent hypothesis at tree depth (j, k) , which is a set of (possibly fused) ground tracks constructed from the set of N ground tracks $\{\tau_i^{(m_i)}\}_{i=1}^N$ obtained by transforming a track cluster subset using the propagation paths $\{m_i\}_{i=1}^N$. It is assumed that for those $\tau_i^{(m_i)}$ which are associated and fused to form the tracks in $\bar{\lambda}$, the paths m_i satisfy constraints of the form $m_{i_1} < m_{i_2}$ for all marginally ordered and associated radar tracks τ_{i_1}, τ_{i_2} . If the parent hypothesis probability $\Pr\{\bar{\lambda} | \tilde{Z}^k\}$ with $\tilde{Z}^k = \cup_{i=1}^N \tilde{Z}_i^k$ is known, then the task is to evaluate $\Pr\{\lambda | \tilde{Z}^k \cup \tilde{Z}_i^{k'}\}$, where λ is a path-dependent association hypothesis on a tree node which branches from the $\bar{\lambda}$ node, and where $\tilde{Z}_i^{k'}$ is the new set of measurements associated with a track $\tau_i^{(m_i)}$ at time $k' \geq k$. The case $k' = k$ corresponds with static hypothesis evaluation, where τ_i is the $(j+1)$ th track in the cluster, so that λ is at depth $(j+1, k)$ in the hypothesis tree. The case $k' > k$ describes hypothesis evaluation at the next time step when either new measurements from existing tracks in the cluster are considered, or where a new track joins the cluster and the number of possible hypotheses expands as illustrated in Figure 2. This latter case requires track state prediction using the dynamic process model in the tracking filter, and is not considered explicitly here. A general track association metric for nondeterministic processes derived in Ref. 22 may be exploited to treat the dynamic case in evaluating hypothesis probabilities and fused state estimates.

If $\tau_i^{(m_i)}$ is not associated with any existing (fused) ground track in $\bar{\lambda}$, then $\tau_i^{(m_i)} \in \lambda$ and $\Pr\{\lambda | \tilde{Z}^k \cup \tilde{Z}_i^{k'}\} = C^{-1} \beta_i^{(m_i)}(k') \Pr\{\bar{\lambda} | \tilde{Z}^k\}$ with normalisation constant C . Since $\tau_i^{(m_i)}$ is then independent of other tracks in λ , the

propagation path m_i may be chosen independent of the utilised paths in $\bar{\lambda}$. Alternatively, $\tau_i^{(m_i)}$ is associated with one of the (fused) tracks, say $\bar{\tau}^{(\bar{m})} \in \bar{\lambda}$, where \bar{m} denotes the set of propagation paths utilised in deriving the set of ground tracks which have been fused to give $\bar{\tau}^{(\bar{m})}$. For this association to be possible, it is assumed that $m_i \notin \bar{m}$ and that m_i satisfies any additional path constraints which may be inferred from the CR advice.

Proceeding by analogy with hierarchical multihypothesis multisensor tracking,^{9,19,22} it follows that the probability of the path-dependent track-to-target hypothesis λ , conditional on the available measurements, is given by the recursive expression

$$\Pr\{\lambda | \bar{Z}^k \cup \bar{Z}_i^{k'}\} = C^{-1} \Pr\{\bar{\lambda} | \bar{Z}^k\} \beta_i^{(m_i)}(k') L(\tau^{(\bar{m}, m_i)} | \bar{Z}^k \cup \bar{Z}_i^{k'}), \quad (23)$$

where $L(\tau^{(\bar{m}, m_i)} | \bar{Z}^k \cup \bar{Z}_i^{k'})$ is the likelihood of the new fused ground track $\tau^{(\bar{m}, m_i)} = \bar{\tau}^{(\bar{m})} \cup \tau_i^{(m_i)} \in \lambda$, and is evaluated using

$$L(\tau^{(\bar{m}, m_i)} | \bar{Z}^k \cup \bar{Z}_i^{k'}) = \int dy(k') p(y(k') | \bar{Z}^k \cup \bar{Z}_i^{k'}, \tau^{(\bar{m}, m_i)}). \quad (24)$$

The integrand in (24) is the pdf of the fused ground state, conditional on the track-associated measurements and on the propagation path set $\bar{m} \cup \{m_i\}$, and is given by

$$p(y(k') | \bar{Z}^k \cup \bar{Z}_i^{k'}, \tau^{(\bar{m}, m_i)}) = C^{-1} \frac{p(y(k') | \bar{Z}^k, \bar{\tau}^{(\bar{m})}) p(y(k') | \bar{Z}_i^{k'}, \tau_i^{(m_i)})}{p(y(k') | \bar{Z}^k \cap \bar{Z}_i^{k'}, \bar{\tau}^{(\bar{m})} \cap \tau_i^{(m_i)})}. \quad (25)$$

The denominator in (25) accounts for track dependence arising from common measurements, the effect of which must be removed to avoid double counting. The generalisation of analogous expressions to the nondeterministic dynamic case is discussed in Ref. 22.

An application of this formalism is provided by binary fusion of associated OTHR radar coordinate tracks τ_1 , τ_2 with ground state pdfs given by Gaussian mixtures of the form (16). Assuming the tracks have no common measurements ($\bar{Z}_1^k \cap \bar{Z}_2^k = \emptyset$) so that the denominator in (25) can be absorbed into the normalisation constant, then (25) is the product of two Gaussian distributions, which is itself Gaussian, and the path-dependent likelihood (24) may be evaluated to give

$$L(\tau_1^{(m_1)} \cup \tau_2^{(m_2)} | \bar{Z}_1^k \cup \bar{Z}_2^k) = \left| 2\pi \left(P_{y_1}^{(m_1)} + P_{y_2}^{(m_2)} \right) \right|^{-\frac{1}{2}} \exp \left[-\frac{1}{2} \left(\hat{y}_1^{(m_1)} - \hat{y}_2^{(m_2)} \right)' \left(P_{y_1}^{(m_1)} + P_{y_2}^{(m_2)} \right)^{-1} \left(\hat{y}_1^{(m_1)} - \hat{y}_2^{(m_2)} \right) \right]. \quad (26)$$

Assuming that there are M propagation modes and that the constraint $m_2 < m_1$ applies to the propagation paths, then the pdf for the fused track ground state is given by the Gaussian mixture

$$p(y(k) | \bar{Z}_1^k \cup \bar{Z}_2^k, \tau_1 \cup \tau_2) = C^{-1} \sum_{m_1=1}^M \sum_{m_2=1}^{m_1-1} \beta_1^{(m_1)} \beta_2^{(m_2)} L(\tau_1^{(m_1)} \cup \tau_2^{(m_2)} | \bar{Z}_1^k \cup \bar{Z}_2^k) \mathcal{N}(y(k); \hat{y}_{1,2}^{(m_1, m_2)}, P_{y_1 y_2}^{(m_1, m_2)}), \quad (27)$$

with the path-dependent fused ground state estimates given by

$$\hat{y}_{1,2}^{(m_1, m_2)} = P_{y_2}^{(m_2)} \left(P_{y_1}^{(m_1)} + P_{y_2}^{(m_2)} \right)^{-1} \hat{y}_1^{(m_1)} + P_{y_1}^{(m_1)} \left(P_{y_1}^{(m_1)} + P_{y_2}^{(m_2)} \right)^{-1} \hat{y}_2^{(m_2)} \quad (28)$$

with covariances

$$P_{y_1 y_2}^{(m_1, m_2)} = P_{y_1}^{(m_1)} \left(P_{y_1}^{(m_1)} + P_{y_2}^{(m_2)} \right)^{-1} P_{y_2}^{(m_2)}. \quad (29)$$

A simple one-dimensional example of binary track fusion when $M = 3$ is illustrated in Figure 4.

4. SUMMARY

An algorithm for the association and fusion of OTHR tracks given by existing tracking filters has been outlined here, based on the evaluation of multiple association hypotheses analogous with multihypothesis multisensor track fusion algorithms reported in the literature.^{9,19} The algorithm details necessary for implementation depend on the tracking filter employed and the form of the CR advice available. Issues such as track dependence due to shared detections, specific physical constraints on propagation path utilisation, false tracks and the association of track segments have yet to be addressed. These extensions to the present work can be accommodated within the formalism presented here.




Figure 4. Example mixture ground state pdfs $p(y(k)|\tilde{Z}_1^k, \tau_1)$ (dashed curve) and $p(y(k)|\tilde{Z}_2^k, \tau_2)$ (dotted curve) for the case of $M = 3$ propagating modes. If the tracks τ_1 and τ_2 are associated, then the mixed fused state pdf is given by $p(y(k)|\tilde{Z}_1^k \cup \tilde{Z}_2^k, \tau)$ with $\tau = \tau_1 \cup \tau_2$ (solid curve). The axis units are arbitrary.

ACKNOWLEDGEMENTS

The authors wish to acknowledge their colleagues in the High Frequency Radar Division for many helpful discussions related to this work.

REFERENCES

1. T. A. Croft, Sky-wave backscatter: a means for observing our environment at great distances, *Rev. Geophys.* **10**, 73, 1972.
2. J. M. Headrick, Looking over the horizon, *IEEE Spectrum*, pp. 36-39, July 1990.
3. McNamara L. F., *The Ionosphere: Communications, Surveillance and Direction finding*, Kreiger Publishing Company, Malabar, Florida, 1991.
4. Y. Bar-Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, New York: YBS, 1995.
5. S. B. Colegrove and P. J. Edwards, Review of the automatic tracking of targets in clutter, *Proc. 21st Institute of Radio and Electronics Engineers IRECON '87 International Electronics Convention and Exhibition*, pp. 681-684, 1987.
6. E. Skafidas and D. J. Percival, Estimation of ionospheric state using a networked array of ionospheric sounders, *Proc. 4th International Symposium on Signal Processing Applications*, pp. 845-848, 1996.
7. E. Skafidas, D. J. Percival and R. J. Evans, Robust ionospheric density estimation using a vertical incidence sounder array, submitted to *IEEE Trans. on Aerospace and Electronic Systems*, May 1997.
8. N. S. Wheadon, J. C. Whitehouse, J. D. Milsom and R. N. Herring, Ionospheric modelling and target coordinate registration for HF sky-wave radars, *HF Radio Systems and Techniques*, IEE Conference Publication **392**, pp. 258-266, July 1994.
9. M. E. Liggins, C. Y. Chong, I. Kadar, M. G. Alford, V. Vannicola and S. Thomopoulos, Distributed fusion architectures and algorithms for target tracking, *Proc. of the IEEE* **85**, 1, pp. 95-107, 1997.
10. I. W. Dall and D. J. Kewley, Track association in the presence of multi-mode propagation, *Proc. International Conference on Radar 92*, IEEE Publication **365**, pp. 70-73, 1992.
11. K. A. B. White, I. W. Dall and A. J. Shellshear, Association of over-the-horizon radar tracks with tracks from microwave radar and other sources, *Proc. SPIE* **2755**, pp. 335-346, 1996.

12. J. Zhu, R. E. Bogner, A. Bouzerdoum and M. L. Southcott, Application of neural networks to track association in over the horizon radar, *Proc. SPIE* **2233**, pp. 224-235, 1994.
13. M. L. Southcott and R. E. Bogner, An association metric for track association of multimode OTHR tracks, submitted to *IEEE Trans. on Aerospace and Electronic Systems*, 1996.
14. G. W. Pulford and R. J. Evans, A multipath data association tracker for over-the-horizon radar, submitted to *IEEE Trans. on Aerospace and Electronic Systems*, December 1996.
15. G. W. Pulford, R. J. Evans, Probabilistic data association for systems with multiple simultaneous measurements, *Automatica* **32**(9), pp. 1311-1316, 1996.
16. T. Kurien, D. Logan and W. P. Berry, Fusion of OTH radar data, *Proc. Seventh Joint Service Data Fusion Symposium*, pp. 221-235, 1994.
17. W. J. Yssel, W. C. Torrez and R. Lematta, Multiple relocatable over-the-horizon radar (ROTHR) track data fusion (MRTDF), *Proc. 9th Nat. Symp. Sensor Fusion*, pp. 37-45, 1996.
18. J. L. Krolik and R. H. Anderson, Maximum likelihood coordinate registration for over-the-horizon radar, accepted, *IEEE Transactions on Signal Processing*.
19. C. Y. Chong, S. Mori, and K. C. Chang, 'Distributed multitarget multisensor tracking', in *Multitarget Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom, Ed., Norwood MA: Artech House, 1990, pp. 247-295.
20. D. B. Reid, An algorithm for tracking multiple targets, *IEEE Trans. Automatic Control*, **AC-24**, pp. 843-854, 1979.
21. L. Comtet, *Advanced Combinatorics - The Art of Finite and Infinite Expansions*, D. Reidel Publishing Company, Dordrecht, Holland, 1974.
22. S. Mori, K. A. Demetri, W. H. Barker and R. N. Lineback, A theoretical foundation of data fusion - generic track association metric, *Proc. Seventh Joint Service Data Fusion Symposium*, pp. 585-594, 1994.
23. M. Abramowitz and I. A. Stegun (Eds.), *Handbook of Mathematical Functions*, Dover Publications, Inc., New York, 1972.

APPENDIX A. Enumeration of Association Hypotheses

The *Stirling number of the second kind*, denoted here by $S(J, T)$, gives the number of track-to-target association hypotheses which assign J distinct tracks to T targets in any order, such that each target has at least one track assigned to it. The sum

$$H(J) = \sum_{T=1}^J S(J, T) \quad (30)$$

then gives the total number of hypotheses when there are J conceivably related tracks. $H(J)$ is called the *exponential number* or the *Bell number* in the literature.²¹ Identities and recurrence relations useful for computing the number of track association hypotheses are summarised here; further discussion may be found in Ref. 21, and additional identities are listed in Ref. 23.

The Stirling number of the second kind $S(J, T)$ is given by the closed form

$$S(J, T) = \frac{1}{T!} \sum_{t=1}^T (-1)^{T-t} {}^T C_t t^J, \quad 1 \leq T \leq J, \quad (31)$$

with

$${}^T C_t = \frac{T!}{(T-t)! t!}. \quad (32)$$

It follows that $S(J, 1) = 1$, $S(J, 2) = 2^{J-1} - 1$, and $S(J, 3) = (3^{J-1} + 1)/2 - 2^{J-1}$. Special values are $S(J, 1) = S(J, J) = 1$ and $S(J, J-1) = {}^J C_2$. If $T > J$ then $S(J, T) = 0$ by definition, which may be interpreted here as there being no hypothesis which assigns J tracks to T targets with no target unassigned when the number of targets exceeds the number of tracks.

It is straightforward to compute the $S(J, T)$ using the triangular recurrence relation

$$S(J, T) = S(J-1, T-1) + TS(J-1, T), \quad J, T \geq 1, \quad (33)$$

J	T										$H(J)$
	1	2	3	4	5	6	7	8	9	10	
1	1	0	0	0	0	0	0	0	0	0	1
2	1	1	0	0	0	0	0	0	0	0	2
3	1	3	1	0	0	0	0	0	0	0	5
4	1	7	6	1	0	0	0	0	0	0	15
5	1	15	25	10	1	0	0	0	0	0	52
6	1	31	90	65	15	1	0	0	0	0	203
7	1	63	301	350	140	21	1	0	0	0	877
8	1	127	966	1701	1050	266	28	1	0	0	4140
9	1	255	3025	7770	6951	2646	462	36	1	0	21147
10	1	511	9330	34105	42525	22827	5880	750	45	1	115975

Table 1. Tabulated values of $S(J, T)$ and $H(J)$ for $1 \leq J \leq 10$ and $1 \leq T \leq 10$.

with $S(J, 0) = S(0, T) = 0$, except $S(0, 0) = 1$. As a check on computation, the identity

$$\sum_{T=0}^J (-1)^{J-T} T! S(J, T) = 1 \quad (34)$$

is useful.

The exponential number $H(J)$ satisfies the recurrence relation

$$H(J+1) = \sum_{j=0}^J {}^J C_j H(j), \quad J \geq 0, \quad (35)$$

and can also be given by the convergent series

$$H(J) = \frac{1}{e} \sum_{j=0}^{\infty} \frac{j^J}{j!} = \left\lfloor \frac{1}{e} \sum_{j=0}^{2J} \frac{j^J}{j!} \right\rfloor, \quad J \geq 1, \quad (36)$$

where the notation $\lfloor \cdot \rfloor$ denotes the nearest integer to the argument.

Representative values of the Stirling numbers of the second kind and the exponential numbers are listed in Table 1.

Multipath Coordinate Registration and Track Fusion for Over-the-horizon Radar

D.J. Percival¹ and K.A.B. White

*Wide Area Surveillance Division,
Defence Science and Technology Organisation,
P.O. Box 1500, Salisbury, S.A. 5108, AUSTRALIA.*

Abstract

Multiple tracks for a single target can be produced by tracking filters used in sky-wave over-the-horizon radar (OTHR) surveillance when several signal propagation paths exist between the target and the radar sites. A multihypothesis multipath track fusion algorithm for OTHR has been described recently which may be appended to existing OTHR tracking filters, and which utilises propagation advice from ionospheric modelling and observations to transform radar tracks to ground coordinates, where track-to-target association hypotheses may be evaluated (D.J. Percival and K.A.B. White, *Proc. SPIE 3163*, to appear, September 1997). Important aspects of this algorithm are developed here, including an efficient strategy for eliminating hypotheses which are physically prohibited, and the explicit enumeration and evaluation of the remaining plausible propagation path-dependent target-to-track association hypotheses.

1 Introduction

Skywave over-the-horizon radar (OTHR) exploits the refraction of HF signals by the ionosphere to detect and track targets beyond the line-of-sight horizon [1]. Several propagation paths between the target and the radar sites may exist in a complex propagation environment, giving multiple resolved detections for a single target. Conventional tracking filters [2] implemented in radar coordinates (e.g., *group range* or *slant range* derived from the signal time delay between transmission and reception, and *apparent azimuth* inferred from the signal reception direction) will then yield multiple radar tracks for each target when multipath propagation conditions are prevalent.

¹ Corresponding Author: John Percival. Email: john.percival@dsto.defence.gov.au

The association and fusion of multipath OTHR tracks has been the subject of several research efforts reported in the literature [3]–[6]. More recently, a multipath probabilistic data association (MPDA) filter has been proposed which directly tracks targets in ground coordinates using multipath detections [7]. Of central importance in these algorithms is the *coordinate registration* (CR) of radar detections or tracks to ground coordinates (e.g., target *ground range* and *true azimuth*), which requires knowledge of the available signal propagation paths [8]. This CR advice can be provided by numerically tracing rays through ionospheric models which are estimated using real-time sounding observations of the ionospheric state. However, even when all propagation paths to a target for a given ionospheric state are well-characterised, the correct assignment of paths to radar tracks is not known *a priori*, and there is some ambiguity in the association of radar tracks with targets when multipath propagation conditions exist.

In [9], a multihypothesis track fusion algorithm for OTHR multipath tracks was outlined, in which all track-to-target association hypotheses were constructed, and the probability of each hypothesis and the fused target states within each hypothesis were recursively evaluated. Important aspects of this algorithm are developed in this brief paper, including a strategy for the exclusion of physically prohibited association hypotheses and the enumeration of the remaining plausible path-dependent target-to-track association hypotheses.

In Section 2, a formalism which incorporates OTHR tracking and CR is summarised, and the transformation of radar track states to ground coordinates using uncertain CR advice is outlined. In Section 3, the formation and enumeration of path-dependent track-to-target association hypotheses is described, and a strategy for the elimination of physically prohibited path assignments is outlined, based on the *marginal ordering* of transformed track states. For a simple illustrative multipath track fusion scenario, the evaluation of each association hypothesis probability is described, and the target ground states arising from the fusion of associated tracks within each hypothesis are estimated.

2 Uncertain Coordinate Registration of OTHR Tracks

OTHR tracking and coordinate registration are jointly considered here as concurrent estimation tasks. Track states in radar coordinates are estimated from associated detections in a multitarget cluttered environment using a conventional tracking filter [2]. Let $x_j(k) \in \mathcal{R}$ be the state at time k associated with the j th track τ_j in the radar coordinate space \mathcal{R} . This state is estimated in a multitarget, cluttered environment from a set of radar measurements

$Z_j^k = \{z_j(i)\}_{i=1}^k$ associated with τ_j using an established tracking filter [2]. The details of track initiation, target dynamic modelling and clutter modelling are assumed to have been addressed in the tracking filter. For Kalman-based filters, it is assumed that the state probability density function (pdf) conditioned on the measurement sequence is well approximated by a Gaussian distribution, denoted here by $\mathcal{N}(x_j(k); \hat{x}_j(k|k), P_{x_j}(k|k))$, with conditional mean $\hat{x}_j(k|k) = E\{x_j(k)|Z_j^k\}$ and covariance $P_{x_j}(k|k) = \text{cov}\{x_j(k)|Z_j^k\}$. The output of the tracking filter is a set of J tracks $\{\tau_j\}_{j=1}^J$; this set forms a cluster of conceivably associated tracks, where two or more tracks may have arisen from detections of a single target propagating via different paths.

Let $w(k) \in \mathcal{I}$ describe the state of the ionosphere at time k in an ionospheric state space \mathcal{I} . The components of $w(k)$ are parameters in a spatial model of the ionospheric refractive index which determines the HF propagation paths between a target and the radar sites. It is assumed that physical arguments and historical observations are sufficient to establish an ionospheric model which is adequate for the coordinate registration of OTHR tracks. Again, Kalman filter-based algorithms may be employed to recursively compute the conditional ionospheric state estimate $\hat{w}(k|k) = E\{w(k)|D^k\}$ and covariance $P_w(k|k) = \text{cov}\{w(k)|D^k\}$ from ionospheric sounding measurements $\{d(i)\}_{i=1}^k$ and prior physical information $d(0)$, collectively denoted here by $D^k = \{d(0), \{d(i)\}_{i=1}^k\}$. An extended Kalman filter has been used to estimate parameters in a simple single-layer ionospheric model from vertical-incidence ionospheric sounding measurements recorded by a geographically distributed array of ionosondes [10]. Multiple-layer ionospheric models are required to describe multipath propagation conditions.

Numerical ray tracing through the ionosphere $w(k) \in \mathcal{I}$ is used to identify the CR transformation $T^{(m)} : \mathcal{R} \times \mathcal{I} \rightarrow \mathcal{G}$ from radar coordinates \mathcal{R} to ground coordinates \mathcal{G} for each propagation path $m \in \mathcal{M}$, where \mathcal{M} is a set of integers which label the available propagation paths. That is, for the state $x_j(k) \in \mathcal{R}$ associated with the radar track τ_j , the corresponding multipath states in ground coordinates are given by

$$y_j^{(m)}(k) = T^{(m)}(x_j(k), w(k)) \in \mathcal{G}, \quad (1)$$

with $m \in \mathcal{M}(x_j(k), w(k))$. Without loss of generality, it may be assumed that $m \in \{1, 2, \dots, M\}$ for all $x_j(k)$ and $w(k)$, where M is the total number of different paths which may exist at any time. Since $x_j(k)$ and $w(k)$ are independent, it is straightforward to show by linearising $T^{(m)}$ about the current state estimates $\hat{x}_j(k|k)$ and $\hat{w}(k|k)$, that the path-dependent ground state conditional estimates and covariances are respectively given by

$$\hat{y}_j^{(m)}(k|k) = T^{(m)}(\hat{x}_j(k|k), \hat{w}(k|k)), \quad (2)$$

with $m \in \mathcal{M}(\hat{x}_j(k|k), \hat{w}(k|k))$, and

$$P_{y_j}^{(m)}(k|k) = J_{x_j}^{(m)} P_{x_j}(k|k) J_{x_j}^{(m)'} + J_w^{(m)} P_w(k|k) J_w^{(m)'}, \quad (3)$$

where the transformation Jacobians $J_{x_j}^{(m)} = \partial T^{(m)} / \partial x_j$ and $J_w^{(m)} = \partial T^{(m)} / \partial w$ are evaluated at $(\hat{x}_j(k|k), \hat{w}(k|k))$, and where the prime denotes the matrix transpose. From (3), it follows that the ground state covariance is the sum of two terms, the first of which is due to the track state uncertainty transformed to ground coordinates, and the second due to the effect of an uncertain ionosphere on the CR transformation. Computational issues which arise in the implementation of this formalism are briefly discussed in [9].

If $\theta_j^{(m)}$ denotes the event that the radar track τ_j is propagating via the m th path, then the conditional probability

$$\beta_j^{(m)}(k) = \Pr\{\theta_j^{(m)} | \tilde{Z}_j^k, \tau_j\} \quad (4)$$

may be defined, with normalisation $\sum_{m=1}^M \beta_j^{(m)}(k) = 1$ assumed, and where, for the sake of brevity, \tilde{Z}_j^k collectively denotes the measurement sequences Z_j^k and D^k . It is anticipated here that the $\beta_j^{(m)}$ can be inferred from the CR advice for each radar track τ_j and for each propagation path $m \in \mathcal{M}$. Since the paths are mutually exclusive, it follows from total probability that the pdf for the j th ground track state $y_j(k) \in \mathcal{G}$ conditioned on the radar and ionospheric measurement sequences is given by the Gaussian mixture

$$p(y_j(k) | \tilde{Z}_j^k, \tau_j) = \sum_{m=1}^M \beta_j^{(m)}(k) \mathcal{N}(y_j(k); \hat{y}_j^{(m)}(k|k), P_{y_j}^{(m)}(k|k)). \quad (5)$$

3 Association Hypothesis Formation and Evaluation

When multipath propagation conditions exist, there may be many plausible path-dependent track-to-target association hypotheses, each resulting in a different distribution of targets in ground coordinates. Track-to-target association hypotheses can be constructed using a *hypothesis tree*, where each route through the tree corresponds with a unique association of radar tracks with targets for a mutually feasible set of available propagation paths. Hypothesis trees for OTHR multipath track fusion were introduced in [9]; the extension of the tree construction to explicitly account for path-dependence is outlined here for the case of static hypothesis evaluation.

For a given cluster of J radar tracks, a hypothesis tree is constructed by taking each track in turn, and associating it with a particular target and propagation




Fig. 1. (a) Path-dependent and (b) path-independent hypothesis trees for a static cluster of three radar tracks when two propagation paths are available. The trees are constructed from left to right as each track labelled $j = 1, 2, 3$ is considered in turn. In (a), the nodes are labelled with the ordered pair (t_j, m_j) , where $t_j = 1, 2, 3$ is a target label and $m_j = 1, 2$ labels the propagation path. In the path-independent case (b), nodes are labelled with $t_j = 1, 2, 3$ only.

path. A static path-dependent hypothesis tree for a cluster of three tracks ($J = 3$), when two propagation paths ($M = 2$) are available is shown in Figure 1(a). Cluster nodes are labelled with the ordered pair (t_j, m_j) , where $t_j = 1, 2, 3$ denotes the target label assigned to the j th radar track, and where $m_j = 1, 2$ labels the path assignment. The first track in the cluster ($j = 1$) is assigned the target label $t_1 = 1$, and may be propagating via either one of two paths ($m_1 = 1$ or $m_1 = 2$). When the second track ($j = 2$) is considered, five hypotheses arise, namely the hypothesis that the second track is associated with the same track as the first ($t_2 = 1$ and $m_2 = 2$ since $m_2 \neq m_1$), and the four hypotheses that the second track is due to a different target ($t_2 = 2$ and $m_2 = 1$ or $m_2 = 2$ for each $j = 1$ hypothesis). Note that the hypothesis where the first two tracks are associated with the same target, but with $m_1 = 2$ and $m_2 = 1$, is excluded by the additional path constraints discussed below. Fourteen hypotheses are generated when the third track ($j = 3$) is considered. The corresponding path-independent track-to-target hypothesis tree is shown in Figure 1(b), and may be derived from the path-dependent tree shown in Figure 1(a) by summing over available propagation paths. Note that the case where all three tracks are associated with the same target ($t_1 = t_2 = t_3 = 1$) is prohibited in this example, since tracks associated with the same target must be propagating via different paths, and only two propagation paths are available here.




Fig. 2. Sketch of three ($J = 3$) radar state estimates and their transformation to ground coordinates via two ($M = 2$) propagation paths, with marginal ordering of states in radar group range and ground range respectively.

Hypothesis formation using an arbitrary ordering of cluster tracks will only permit a path constraint of the form $m_i \neq m_j$ to be imposed for the association of two radar track states \hat{x}_i and \hat{x}_j with the same target. However, if cluster tracks are *marginally ordered*, other physically prohibited path combinations can be eliminated *a priori*. A set of track state estimates $\{\hat{x}_j\}_{j=1}^J$ is marginally ordered with respect to a given state vector component if it is arranged such that the prescribed component is ordered with increasing magnitude. For example, when M propagation paths are available, the $J \leq M$ radar state estimates associated with a given target may be marginally ordered with respect to *increasing* group range, while for given estimates \hat{x}_j and \hat{w} , the propagation path labels m_j may be ordered such that the transformed ground states $\hat{y}_j^{(m_j)}$, $m_j \in \{1, 2, \dots, M\}$ are marginally ordered with *decreasing* ground range. If track states are then considered in this marginal order when forming the hypothesis tree, the assignment of the path $m_j \in \{1, 2, \dots, M\}$ to the j state estimate \hat{x}_j will physically constrain the path assignment subsequently made to \hat{x}_{j+1} to be an element of the set $m_{j+1} \in \{m_j + 1, m_j + 2, \dots, M\}$, for those hypotheses which associate \hat{x}_j and \hat{x}_{j+1} with the same target. In Figure 2, the CR transformation and marginal ordering of state estimates is shown schematically for the scenario corresponding with Figure 1, where three radar track states are transformed using two propagation paths.

The number of track-to-target association hypotheses for the path-independent case was evaluated in [9] as a function $H(J)$ which could be expressed in terms of *Stirling numbers of the second kind* [11]. When propagation path availability is also considered, the number of track-to-target association hypotheses which may be formed greatly increases. If there are J tracks and M propagation paths, it may be shown as an exercise in combinatorics that the number $\tilde{H}(J, M)$ of path-dependent track-to-target association hypotheses is given by

J	M								$H(J)$
	1	2	3	4	5	6	7	8	
1	1	2	3	4	5	6	7	8	1
2	1	5	12	22	35	51	70	92	2
3	1	14	55	140	285	506	819	1240	5
4	1	43	282	1005	2630	5706	10913	19062	15
5	1	142	1578	7924	26751	71082	160888	324584	52
6	1	499	9514	67666	295530	963473	2583973	6026420	203

Table 1

Tabulated values of $\tilde{H}(J, M)$ and $H(J)$ for $1 \leq J \leq 6$ and $1 \leq M \leq 8$.

the recurrence relation

$$\tilde{H}(J, M) = \sum_{j=\max(0, J-M)}^{J-1} \binom{J-1}{j} \binom{M}{J-j} \tilde{H}(j, M) \quad (6)$$

for $J \geq 1$, and with $\tilde{H}(0, M) = 1$ for $M \geq 1$. It does not appear that $\tilde{H}(J, M)$ has been discussed in the literature; $H(J)$ and $\tilde{H}(J, M)$ are evaluated and compared in Table 1 for $1 \leq J \leq 6$ and $1 \leq M \leq 8$. Note that the cases $\tilde{H}(1, 2) = 1$, $\tilde{H}(2, 2) = 5$ and $\tilde{H}(3, 2) = 14$ are illustrated in Figure 1(a).

A recursive algorithm was outlined in [9] to compute the probability of each path-dependent track-to-target association hypothesis, given the probability of the parent hypothesis on the preceding node along the same branch of the hypothesis tree. At the tree root node when the first track τ_1 is considered, there are M path-dependent association hypotheses, denoted $\lambda_1^{m_1}$ for $m_1 = 1, 2, \dots, M$, each with probability

$$\Pr\{\lambda_1^{m_1} | \tilde{Z}_1^k\} = \Pr\{\tau_1 \cap \theta_1^{(m_1)} | \tilde{Z}_1^k\} = \beta_1^{(m_1)}(k). \quad (7)$$

If $\lambda_{t_1 t_2 \dots t_j}^{m_1 m_2 \dots m_j}$ denotes the path-dependent hypothesis at time k in which the marginally ordered radar tracks $\tau_1, \tau_2, \dots, \tau_j$ are taken as propagating via the paths m_1, m_2, \dots, m_j respectively and are associated with targets labelled t_1, t_2, \dots, t_j respectively, a general expression for evaluating $\Pr\{\lambda_{t_1 t_2 \dots t_{j+1}}^{m_1 m_2 \dots m_{j+1}} | \cup_{i=1}^{j+1} \tilde{Z}_i^k\}$ from its parent hypothesis probability $\Pr\{\lambda_{t_1 t_2 \dots t_j}^{m_1 m_2 \dots m_j} | \cup_{i=1}^j \tilde{Z}_i^k\}$ is given in [9]. This formulation requires the evaluation of a track association likelihood for those hypotheses where two or more ground tracks are associated with the same target. For the sake of brevity, this general formulation is not repeated here. However, for the scenario with three radar tracks and two propagation paths illustrated in Figure 2, there are four path-independent hypotheses as shown in Figure 1(b), which are denoted $\lambda_{t_1 t_2 t_3}$ with the relevant target label




Fig. 3. Example target ground state pdfs for the four path-independent association hypotheses generated for three radar tracks and two propagation paths. Individual target pdfs are distinguished by line style; the units are arbitrary here.

sequence given by the node labels along each branch of the hypothesis tree. Example target pdfs as a function of ground range for each hypothesis are illustrated in Figure 3. For the hypothesis λ_{123} , there are three independent targets, each with a two-component Gaussian mixture pdf arising from the two propagation paths. For the remaining hypotheses λ_{112} , λ_{121} and λ_{122} , there are two target pdfs, one of which is a Gaussian distribution arising from the fusion of the relevant path components of the individual ground state mixture pdfs.

4 Summary

In this brief paper, several developments of a recently described multihypothesis multipath OTHR track fusion algorithm are presented and illustrated with a simple track fusion example. In particular, a technique for the formation and enumeration of the path-dependent track-to-target association hypotheses is given, and a method for imposing physical path constraints on the association hypotheses is described. Further developments of this algorithm, including the treatment of non-deterministic target dynamics in evaluating the association hypotheses, may follow by exploiting the partial analogy with hierarchical multisensor track fusion algorithms reported in the literature [12].

References

- [1] T. A. Croft, Sky-wave backscatter: a means for observing our environment at great distances, *Rev. Geophys.* **10**, 73, 1972.
- [2] Y. Bar-Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, New York: YBS, 1995.
- [3] I. W. Dall and D. J. Kewley, Track association in the presence of multi-mode propagation, *Proc. International Conference on Radar 92*, IEEE Publication **365**, pp. 70-73, 1992.
- [4] J. Zhu, R. E. Bogner, A. Bouzerdoun and M. L. Southcott, Application of neural networks to track association in over the horizon radar, *Proc. SPIE* **2233**, pp. 224-235, 1994.
- [5] T. Kurien, D. Logan and W. P. Berry, Fusion of OTH radar data, *Proc. Seventh Joint Service Data Fusion Symposium*, pp. 221-235, 1994.
- [6] J. L. Krolik and R. H. Anderson, Maximum likelihood coordinate registration for over-the-horizon radar, *IEEE Trans. on Signal Proc.* **45**, 935-959, 1997.
- [7] G. W. Pulford and R. J. Evans, A multipath data association tracker for over-the-horizon radar, submitted to *IEEE Trans. on Aerospace and Electronic Systems*, December 1996.
- [8] N. S. Wheadon, J. C. Whitehouse, J. D. Milsom and R. N. Herring, Ionospheric modelling and target coordinate registration for HF sky-wave radars, *HF Radio Systems and Techniques*, IEE Conference Publication **392**, pp. 258-266, July 1994.
- [9] D. J. Percival and K. A. B. White, Multipath Track Fusion for Over-the-horizon Radar, *Proc. SPIE* **3163**, to appear, September 1997.
- [10] E. Skafidas and D. J. Percival, Estimation of ionospheric state using a networked array of ionospheric sounders, *Proc. 4th International Symposium on Signal Processing Applications*, pp. 845-848, 1996.
- [11] M. Abramowitz and I. A. Stegun (Eds.), *Handbook of Mathematical Functions*, Dover Publications, Inc., New York, 1972.
- [12] M. E. Liggins, C. Y. Chong, I. Kadar, M. G. Alford, V. Vannicola and S. Thomopoulos, Distributed fusion architectures and algorithms for target tracking, *Proc. of the IEEE* **85**, 1, pp. 95-107, 1997.

Multihypothesis Fusion of Multipath Over-the-horizon Radar Tracks

D.J. Percival and K.A.B. White,
Wide Area Surveillance Division,
Defence Science and Technology Organisation,
P.O. Box 1500, Salisbury, S.A. 5108, AUSTRALIA.

ABSTRACT

Multiple tracks for a single target arise in skywave over-the-horizon radar (OTHR) tracking, due to multiple ionospheric propagation paths between the target and the radar sites. A multihypothesis track fusion algorithm was developed and reported in an earlier paper, where all possible path-dependent track-to-target association hypotheses were constructed, and the probability of each hypothesis evaluated [D.J. Percival and K.A.B. White *Proc SPIE* 3163, pp. 363-374, 1997]. The implementation of recursive hypothesis evaluation and fused track state estimation is the subject of this paper. Sources of multipath track dependence are identified, and their treatment discussed. The time evolution of the track-to-target hypothesis probabilities and target state estimates is illustrated for example multitarget OTHR tracking scenarios using a simple stochastic ionospheric model which admits multipath propagation.

Keywords: multipath track fusion, multihypothesis tracking, over-the-horizon radar

1. INTRODUCTION

Skywave over-the-horizon radar (OTHR) exploits ionospheric propagation for wide area surveillance of air and sea targets beyond the line-of-sight horizon.^{1,2} Separate refractive layers in the ionosphere may support multiple propagation paths between targets and the radar site, so that several resolved OTHR detections of single target may be made. Tracking filters used in existing OTHR systems then produce multiple tracks for each target. For OTHR surveillance tasks in which several radar tracks arise from an unknown number of targets, the challenge is to correctly associate radar tracks with targets and then to estimate target locations in ground coordinates by fusing associated track data. Unfortunately, several plausible track-to-target association hypotheses usually exist, and this association ambiguity must be explicitly considered.

Several research efforts have examined the association and fusion of multipath OTHR data.³⁻⁸ The most widely accepted approach is to first track targets in radar coordinates, and then use independent coordinate registration (CR) advice, derived from ionospheric models and radio sounding data, to transform radar tracks to ground coordinates, after identifying the propagation path for each track. An alternative approach which directly tracks targets in ground coordinates using multipath detections has also been reported.⁹

In Ref. 10, a multipath track fusion algorithm was described in which all track-to-target association hypotheses were constructed using all plausible path combinations, and the probability of each hypothesis evaluated. For each hypothesis, target states in ground coordinates were described by a Gaussian mixture probability density function (pdf), taking account of both the track covariance and the uncertainty in the CR advice provided. The CR uncertainty has two parts, namely a discrete part corresponding to uncertainty in the propagation path assignment, and a continuous part arising from the CR transformation uncertainty for a given propagation path due to incomplete knowledge of the ionospheric state. An extension of the algorithm provided a strategy for excluding physically prohibited path-dependent track-to-target association hypotheses by exploiting the marginal order of radar tracks and propagation paths.¹¹

In this paper, the implementation of the multipath track fusion algorithm of Ref. 10 is detailed, with particular attention given to the recursive evaluation of path-dependent track-to-target association hypotheses. In Section 2, OTHR tracking and CR are reviewed, and the formation and evaluation of the association hypotheses for multipath track fusion is outlined. In Section 3, the multipath track fusion algorithm is applied to a simple static track fusion scenario. The generalisation of the algorithm to include scene dynamics is discussed and illustrated in Section 4.

Corresponding Author: John Percival; Email: john.percival@dsto.defence.gov.au

2. MULTIPATH TRACK FUSION ALGORITHM

A multipath track fusion algorithm is outlined here, and the notation required in the sequel is introduced. Important topics such as OTHR tracking and coordinate registration are briefly addressed and the formation and evaluation of path-dependent track-to-target hypotheses is described.

2.1. OTHR Tracking and Coordinate Registration

If $x_j(k) \in \mathcal{R}$ denotes the target state at time k associated with the j th track τ_j in radar coordinates \mathcal{R} , then conventional OTHR tracking filters yield the state pdf $p(x_j(k)|Z_j^k)$ conditional on the sequence $Z_j^k = \{z_j(i)\}_{i=1}^k$ of measurements z_j associated with τ_j . If a Kalman-based tracking filter is used, $p(x_j(k)|Z_j^k) = \mathcal{N}(x_j(k); \hat{x}_j(k|k), P_{x_j}(k|k))$, where \mathcal{N} denotes a Gaussian distribution with conditional mean $\hat{x}_j(k|k) = E\{x_j(k)|Z_j^k\}$ and covariance $P_{x_j}(k|k) = \text{cov}\{x_j(k)|Z_j^k\}$. Typically, in OTHR applications, $x_j(k)$ has four components, namely the *group* or *slant range*, the *apparent azimuth*, the *group range rate* and the *apparent azimuth rate*.

Coordinate registration of radar tracks to ground coordinates requires knowledge of the signal propagation path, usually gained by numerical ray tracing through ionospheric models fitted to radio sounding data. For an estimate of the ionospheric state $w(k) \in \mathcal{I}$ at time k , the CR transformation $T^{(m)} : \mathcal{R} \times \mathcal{I} \rightarrow \mathcal{G}$ to ground coordinates \mathcal{G} may be computed for each propagation path, indexed here by $m \in \{1, 2, \dots, M\}$. Ionospheric state estimation using a sequence of radio sounding measurements $D^k = \{d(i)\}_{i=1}^k$ is the subject of several technical papers.¹²⁻¹⁴ In most circumstances, $T^{(m)}$ can be linearised about the state estimates $\hat{x}_j(k|k)$ and $\hat{w}(k|k) = E\{w(k)|D^k\}$, so that the path-dependent ground state $y_j^{(m)}(k)$, corresponding to the j th radar track propagating via the m th path, is also Gaussian distributed. That is,

$$p(y_j^{(m)}(k)|Z_j^k, D^k) = \mathcal{N}(y_j^{(m)}(k); \hat{y}_j^{(m)}(k|k), P_{y_j}^{(m)}(k|k)), \quad (1)$$

with conditional mean

$$\hat{y}_j^{(m)}(k|k) = T^{(m)}(\hat{x}_j(k|k), \hat{w}(k|k)), \quad (2)$$

and covariance

$$P_{y_j}^{(m)}(k|k) = J_{x_j}^{(m)} P_{x_j}(k|k) J_{x_j}^{(m)'} + J_w^{(m)} P_w(k|k) J_w^{(m)'}, \quad (3)$$

where $P_w(k|k) = \text{cov}\{w(k)|D^k\}$ describes the uncertainty in the ionospheric state estimate. In writing (3), the transformation Jacobians $J_{x_j}^{(m)}(k) = \partial T^{(m)}(x_j, w)/\partial x_j$ and $J_w^{(m)}(k) = \partial T^{(m)}(x_j, w)/\partial w$ are evaluated at $(\hat{x}_j(k|k), \hat{w}(k|k))$, and the prime denotes the matrix transpose. From total probability, and assuming the independence of propagation paths, it follows that the path-independent target ground state distribution is given by the Gaussian mixture

$$p(y_j(k)|Z_j^k, D^k) = \sum_{m=1}^M \beta_j^{(m)}(k) \mathcal{N}(y_j^{(m)}(k); \hat{y}_j^{(m)}(k|k), P_{y_j}^{(m)}(k|k)), \quad (4)$$

where $\beta_j^{(m)}(k) = \Pr\{\theta_j^{(m)}|Z_j^k, D^k, \tau_j\}$ is the probability of the event $\theta_j^{(m)}$ that the radar track τ_j is propagating via the m th path, conditional on the available measurements.

To illustrate this formalism, the CR transformation for the most simple stochastic ionospheric model which admits multipath propagation is described in Appendix A.

2.2. Multihypothesis Multipath Track Fusion

When the surveillance volume contains an unknown number of targets, and where multipath propagation conditions exist, many *path-dependent* track-to-target association hypotheses may be constructed from a cluster of conceivably associated tracks. Each hypothesis corresponds with a distribution of targets in ground coordinates, computed from a unique association of radar tracks with targets for a set of mutually compatible propagation paths. A path-dependent hypothesis for a cluster $\{\tau_j\}_{j=1}^J$ of J radar tracks may be denoted $\lambda_{t_1 t_2 \dots t_J}^{m_1 m_2 \dots m_J}$, where each track τ_j is assigned the path m_j , and is associated with the target t_j . *Path-independent* hypotheses, denoted $\lambda_{t_1 t_2 \dots t_J}$, may be derived from the path-dependent hypotheses by summing over the set of possible path assignments for each target. Path-dependence is only considered in order to apply the correct CR transformation to radar tracks; path-independent hypotheses




Figure 1. Two nodes of a path-dependent hypothesis tree, showing the relationship between a new path-dependent hypothesis λ and its parent hypothesis $\bar{\lambda}$, when the next transformed cluster track $\tau_i^{(m_i)}$ is considered.

focus on the distribution of targets in ground coordinates, the estimation of which is the primary OTHR surveillance mission.

Hypothesis trees have been introduced for the construction of association hypotheses by considering each radar track in a cluster in turn.^{10,11} The extension of the hypothesis tree representation to include target dynamics has also been reported.¹⁰ A key component of the multipath track fusion algorithm is the recursive evaluation of hypothesis probabilities, as the hypothesis tree is constructed. An arbitrary branch of a *path-dependent* hypothesis tree is shown in Figure 1 as a solid line connecting two nodes, where a new path-dependent hypothesis λ is constructed from a parent hypothesis $\bar{\lambda}$, when the next ground track $\tau_i^{(m_i)}$ in the cluster is considered. The notation $\tau_i^{(m_i)}$ used here makes explicit the assignment of the m_i th propagation path to the i th radar track τ_i when applying the CR transformation.

Two possibilities for λ arise: either $\tau_i^{(m_i)}$ is associated with a target already existing in the parent hypothesis $\bar{\lambda}$, or $\tau_i^{(m_i)}$ describes a new target. In the latter case, no multipath track fusion is required, and $\lambda = \bar{\lambda} \cup \{\tau_i^{(m_i)}\}$, with no constraint on the assignment of the propagation path m_i . Otherwise, $\tau_i^{(m_i)}$ is associated with some (possibly fused) target track $\bar{\tau}^{\bar{m}} \in \bar{\lambda}$, where \bar{m} denotes a set of paths which have already been assigned in the construction of $\bar{\lambda}$. The task is then to fuse $\tau_i^{(m_i)}$ with $\bar{\tau}^{\bar{m}}$ to yield the composite target track $\tau^{\bar{m}, m_i} = \bar{\tau}^{\bar{m}} \cup \tau_i^{(m_i)} \in \lambda$. In this case, a physical constraint exists on the assignment of the path m_i . An efficient strategy for excluding prohibited path assignments, based on the marginal ordering of radar tracks, has been described¹¹ which leads to a constraint of the form $m_i > \max(\bar{m})$. The physical explanation for this constraint is straightforward. Associated radar tracks may not arise from the same propagation path, so that $m_i \notin \bar{m}$. Furthermore, for two associated tracks separated in group range, for example, the track with the longer group range cannot be assigned a propagation path with a path length which is shorter than the path assigned to the other track. An extension of this argument to many cluster tracks yields the stated path constraint, where the direction of the inequality is determined by an arbitrary ordering convention. For propagation paths of equal length, the marginal ordering path constraint can be extended to other track state components, such as apparent azimuth.

Expressions for the recursive computation of path-dependent track-to-target association hypothesis probabilities and fused target state estimates were given in Ref. 10, based on an adaptation of existing multisensor track fusion formalism.^{15,16} For the two hypothesis cases illustrated in Figure 1, the probability of the path-dependent track-to-target association hypothesis λ at time k , conditional on all the available measurement data, is given by

$$\Pr\{\lambda | \bar{Z}^k \cup Z_i^k, D^k\} = \begin{cases} C^{-1} \Pr\{\bar{\lambda} | \bar{Z}^k, D^k\} \beta_i^{(m_i)}(k), & \tau_i^{(m_i)} \in \lambda, \tau_i^{(m_i)} \notin \bar{\lambda}, \\ C^{-1} \Pr\{\bar{\lambda} | \bar{Z}^k, D^k\} \beta_i^{(m_i)}(k) L(\tau^{\bar{m}, m_i} | \bar{Z}^k, Z_i^k, D^k), & \tau^{\bar{m}, m_i} = \bar{\tau}^{\bar{m}} \cup \tau_i^{(m_i)} \in \lambda, \end{cases} \quad (5)$$

where $L(\tau^{\bar{m}, m_i} | \bar{Z}^k, Z_i^k, D^k)$ is the likelihood of the fused ground track $\tau^{\bar{m}, m_i}$, with \bar{Z}^k denoting the set of measurement data up to time k associated with $\bar{\tau}^{\bar{m}}$, and where C^{-1} is a normalisation constant. When $\tau_i^{(m_i)}$ is associated with a new target so that $\tau_i^{(m_i)} \in \lambda$, the corresponding path-dependent target pdf in ground coordinates is given by (1) when the index substitutions $j \rightarrow i$ and $m \rightarrow m_i$ are made. The summation over paths then gives the target pdf

for the path-independent hypothesis (4). When $\tau_i^{(m_i)}$ is associated with an existing target $\bar{\tau}^{(\bar{m})} \in \bar{\lambda}$, the fused target pdf, based on the joint information, is given by

$$p(y(k)|\bar{Z}^k \cup Z_i^k, D^k, \tau^{(\bar{m}, m_i)}) = C^{-1} \frac{p(y(k)|\bar{Z}^k, D^k, \bar{\tau}^{(\bar{m})})p(y(k)|Z_i^k, D^k, \tau_i^{(m_i)})}{p(y(k)|\bar{Z}^k \cap Z_i^k, D^k, \bar{\tau}^{(\bar{m})}, \tau_i^{(m_i)})} \quad (6)$$

The denominator in (6) accounts for common information between $\tau_i^{(m_i)}$ and $\bar{\tau}^{(\bar{m})}$, the effect of which must be cancelled to avoid double counting. The substitution of (6) into

$$L(\tau^{(\bar{m}, m_i)}|\bar{Z}^k, Z_i^k, D^k) = \int dy(k) p(y(k)|\bar{Z}^k \cup Z_i^k, D^k, \tau^{(\bar{m}, m_i)}) \quad (7)$$

gives the track association likelihood required for the hypothesis probability computation (5).

For Gaussian distributed pdfs $p(y(k)|\bar{Z}^k, D^k, \bar{\tau}^{(\bar{m})}) = \mathcal{N}(y(k); \hat{y}^{(\bar{m})}(k|k), P_{\bar{y}}^{(\bar{m})}(k|k))$ and $p(y(k)|Z_i^k, D^k, \tau_i^{(m_i)}) = \mathcal{N}(y(k); \hat{y}_i^{(m_i)}(k|k), P_{y_i}^{(m_i)}(k|k))$ in the numerator of (6), the integral (7) can be evaluated to give the path-dependent track association likelihood

$$L(\tau^{(\bar{m}, m_i)}|\bar{Z}^k, Z_i^k, D^k) = \rho^{-1} |2\pi\Pi|^{-\frac{1}{2}} \exp \left[-\frac{1}{2} (\hat{y}^{(\bar{m})} - \hat{y}_i^{(m_i)})' \Pi^{-1} (\hat{y}^{(\bar{m})} - \hat{y}_i^{(m_i)}) \right], \quad (8)$$

with

$$\Pi = P_{\bar{y}}^{(\bar{m})} + P_{y_i}^{(m_i)}, \quad (9)$$

and where ρ is the *a priori* target density. In writing (8) and (9), the arguments $(k|k)$ have been omitted for clarity. The derivation of (8) assumes that there are no common measurements between $\bar{\tau}^{(\bar{m})}$ and $\tau_i^{(m_i)}$, (i.e., $\bar{Z}^k \cap Z_i^k = \emptyset$), and that the denominator in (6) may be approximated by a constant and incorporated into ρ^{-1} . The path-dependent fused ground state pdf is then Gaussian distributed with conditional mean

$$\hat{y}^{(\bar{m}, m_i)} = P_{y_i}^{(m_i)} \Pi^{-1} \hat{y}^{(\bar{m})} + P_{\bar{y}}^{(\bar{m})} \Pi^{-1} \hat{y}_i^{(m_i)}, \quad (10)$$

and covariance

$$P^{(\bar{m}, m_i)} = P_{\bar{y}}^{(\bar{m})} \Pi^{-1} P_{y_i}^{(m_i)}, \quad (11)$$

where the independence of $y_i^{(m_i)}$ and $\bar{y}^{(\bar{m})}$ has again been assumed.

The independence assumption for associated multipath radar tracks requires further consideration. Track dependence due to common process noise has been treated for the fusion of tracks from multiple sensors.^{17,18} The same arguments may be applied to OTHR track fusion to account for common process noise, with the inclusion of a cross-covariance $P_{\bar{y}y_i}^{(\bar{m})(m_i)} = E\{(\bar{y}^{(\bar{m})} - \hat{y}^{(\bar{m})})(y_i^{(m_i)} - \hat{y}_i^{(m_i)})' | \bar{Z}^k \cup Z_i^k, D^k\}$ to describe the track dependence. Specifically, Π in (8) is replaced with

$$\Pi^* = \Pi - P_{\bar{y}y_i}^{(\bar{m})(m_i)} - P_{y_i\bar{y}}^{(m_i)(\bar{m})}, \quad (12)$$

where $P_{\bar{y}y_i}^{(\bar{m})(m_i)} = P_{y_i\bar{y}}^{(m_i)(\bar{m})}$, and where the maximum likelihood fused track conditional mean and covariance are respectively given by¹⁸

$$\hat{y}^{(\bar{m}, m_i)} = \hat{y}^{(\bar{m})} + (P_{\bar{y}}^{(\bar{m})} - P_{\bar{y}y_i}^{(\bar{m})(m_i)}) \Pi^{*-1} (\hat{y}_i^{(m_i)} - \hat{y}^{(\bar{m})}), \quad (13)$$

and

$$P^{(\bar{m}, m_i)} = P_{\bar{y}}^{(\bar{m})} - (P_{\bar{y}}^{(\bar{m})} - P_{\bar{y}y_i}^{(\bar{m})(m_i)}) \Pi^{*-1} (P_{\bar{y}}^{(\bar{m})} - P_{y_i\bar{y}}^{(m_i)(\bar{m})}). \quad (14)$$

The recursive computation of $P_{\bar{y}y_i}^{(\bar{m})(m_i)}$ using the underlying Kalman filter tracking model is described in Ref. 17. It is straightforward to modify this derivation for OTHR multipath track fusion by using the CR transformation (2) to account for the fact that the tracking filter is applied in radar coordinates, while track fusion is performed in ground coordinates. However, the real-time calculation of cross-covariance may not be feasible, so approximate forms are required. An expedient cross covariance approximation¹⁷ is $P_{\bar{y}y_i}^{(\bar{m})(m_i)} \approx \gamma P_{\bar{y}}^{(\bar{m})}$ for the simple case when

$P_{\bar{y}}^{(\bar{m})} \approx P_{y_i}^{(m_i)}$, where γ is chosen such that the ratio of the error ellipsoid volumes $|P_{\bar{y}y_i}^{(\bar{m})(m_i)}|^{1/2} / |P_{\bar{y}}^{(\bar{m})}|^{1/2}$ satisfies the experimentally observed value of 0.7. When the target ground state has four components, $\gamma \approx 0.67$.

Two additional sources of track dependence in OTHR multipath track fusion may be identified. First, it is quite possible that multipath track estimates are updated using common measurements when two or more multipath tracks lie within the measurement gate of the tracking filter. The importance of this effect depends on the details of the measurement association scheme, the radar clutter environment, and the proximity of the multipath detections in radar coordinates. Visibility into the radar tracking algorithm is required to identify common measurements, so that their effect on track association likelihood can be removed by evaluating the denominator in (6). The removal of redundant information is well-studied in the context of distributed estimation,^{15,16} where the *information graph* is shown to be a useful device for tracing information flow. Second, the independence of propagation paths is assumed in the derivation of the target ground state distribution (4). While separate ionospheric layers may well be independent, two different paths can share the same transmission or reception path segment which propagates via a common ionospheric layer. This possible source of track dependence has been neglected here.

3. MULTIPATH HYPOTHESIS EVALUATION: STATIC EXAMPLE

The implementation of the multipath track fusion algorithm is illustrated using a cluster of three radar tracks $\{\tau_1, \tau_2, \tau_3\}$, with four propagation paths available, indexed with $m \in \{1, 2, 3, 4\}$. This scenario is consistent with the propagation model developed in Appendix A, for a two-layer, spherically-symmetric ionosphere. Static hypothesis testing refers to the evaluation of track-to-target association likelihoods and target state estimates at a given time instant, without regard for past track history. The path-independent hypothesis tree for the cluster of three tracks is shown in Figure 2, together with representative target pdfs for each hypothesis. It is assumed throughout this section that the radar track states are independent and Gaussian distributed, with state estimates and covariances given by a tracking filter.

Cluster tracks τ_j are marginally ordered at a given time such that their corresponding group range estimates \hat{R}_j are increasing, i.e., $\hat{R}_1 \leq \hat{R}_2 \leq \hat{R}_3$. Likewise, propagation paths may be ordered with increasing path length, such that for a given τ_j , the transformed ground range estimates $\hat{r}_j^{(m)}$ for each path m satisfy $\hat{r}_j^{(1)} \geq \hat{r}_j^{(2)} \geq \hat{r}_j^{(3)} \geq \hat{r}_j^{(4)}$. Given the uncertainty in track state estimates, it is possible that the ordering of states on the basis of their estimates will be in error. Track ordering can be formulated as a hypothesis test. For example, the ordering of two independent Gaussian random variables $R_1 \sim \mathcal{N}(\hat{R}_1, \sigma_1)$ and $R_2 \sim \mathcal{N}(\hat{R}_2, \sigma_2)$ can be expressed as the hypothesis test

$$\begin{aligned} H_0 : \quad R_2 - R_1 &< 0, \\ H_1 : \quad R_2 - R_1 &\geq 0, \end{aligned}$$

subject to an acceptable "false alarm" probability $\Pr\{H_1|H_0\}$. Hypothesis H_1 for $R_1 \leq R_2$ may be accepted over the null hypothesis H_0 when the test statistic satisfies

$$\frac{\hat{R}_2 - \hat{R}_1}{(\sigma_1^2 + \sigma_2^2)^{1/2}} \geq T_0, \quad (15)$$

with the threshold T_0 chosen using the upper tail of a normal density for a given $\Pr\{H_1|H_0\}$, e.g., $T_0 = 1.64$ for $\Pr\{H_1|H_0\} = 5\%$.

Ordered target tracks in a cluster are considered in turn, and all path-dependent track-to-target association hypotheses are formed with the aid of a hypothesis tree. The first track τ_1 is assigned to target $t_1 = 1$ via one of four paths, yielding the set of hypotheses $\{\lambda_1^1, \lambda_1^2, \lambda_1^3, \lambda_1^4\}$, with $\lambda_1^{m_1} = \{\tau_1^{(m_1)}\}$ for $m_1 \in \{1, 2, 3, 4\}$. The probability of each hypothesis is simply the probability that the path m_1 is utilised for radar track τ_1 , so that

$$\Pr\{\lambda_1^{m_1} | Z_1^k, D^k, \tau_1\} = \beta_1^{(m_1)}(k). \quad (16)$$

The path-dependent target ground state is given by (1) with $j = 1$ and $m = m_1$. Summing over available paths yields the path-independent target ground state pdf, given by the Gaussian mixture (4), with $j = 1$ and $M = 4$. A representative one-dimensional Gaussian mixture pdf is shown in Figure 2 for the hypothesis λ_1 , where three mixture components are evident; for the ionospheric model detailed in Appendix A, the mixed paths $m = 2$ and $m = 3$ have

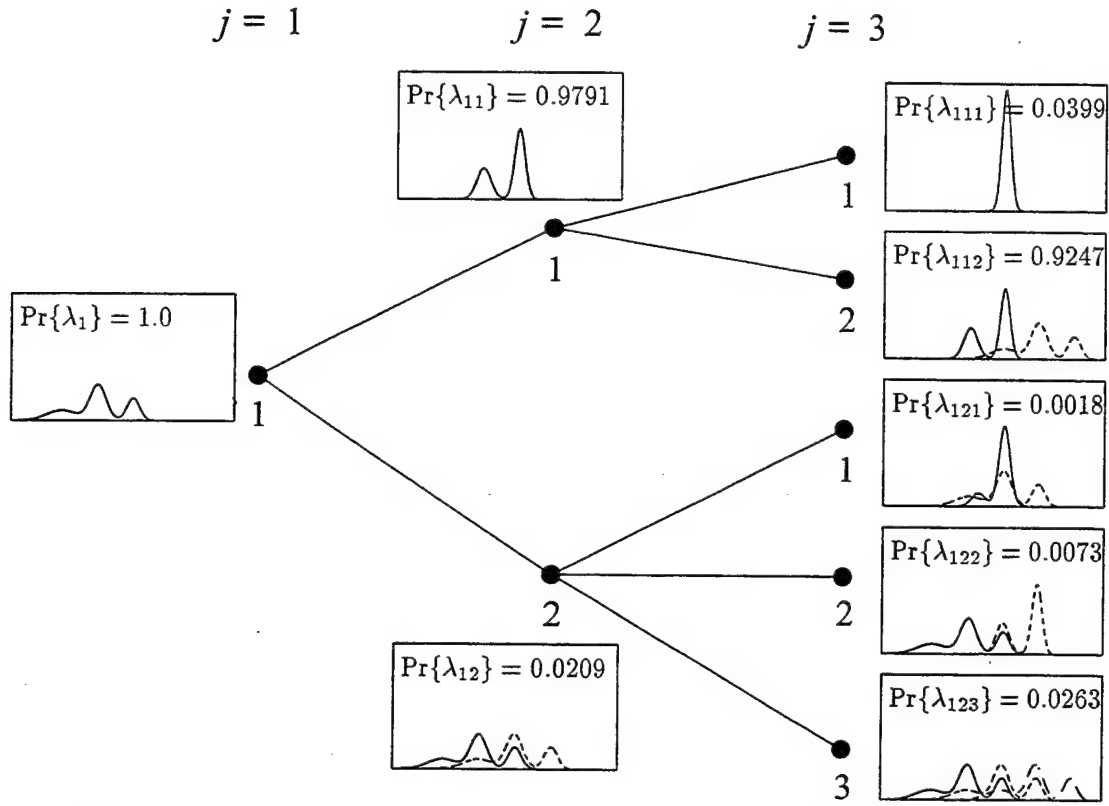


Figure 2. A path-independent hypothesis tree for a cluster of three radar tracks indexed $j \in \{1, 2, 3\}$. Each tree node is labelled with a target index $t_j \in \{1, 2, 3\}$. Representative one-dimensional target pdfs for each path-independent hypothesis are plotted as a function of ground range, with different line types used to distinguish targets. Example hypothesis probabilities are given.

identical path length, and so their contributions to the ground state pdf are not resolved here. When there are no false targets, the probability $\Pr\{\lambda_1|Z_1^k, D^k, \tau_1\}$ of the path-independent hypothesis λ_1 is unity, which follows from the normalisation $\sum_{m=1}^4 \beta_1^{(m)} = 1$. It is straightforward to extend the hypothesis set to consider false tracks by introducing a “false target” $t_j = 0$, to which radar tracks may be assigned.

When the second track τ_2 in the cluster is considered, it is either associated with the same target as the first track ($t_2 = t_1 = 1$) or it is associated with a new target $t_2 = 2$. In the former case, there are six path-dependent hypotheses, labelled $\lambda_{11}^{m_1 m_2} = \{\tau_1^{(m_1)} \cup \tau_2^{(m_2)}\}$, with $m_1, m_2 \in \{1, 2, 3, 4\}$ and subject to the constraint $m_2 > m_1$. In the latter case, the path assignment for τ_2 is unconstrained, and there are $4^2 = 16$ path-dependent hypotheses, labelled $\lambda_{12}^{m_1 m_2} = \{\tau_1^{(m_1)}, \tau_2^{(m_2)}\}$, again with $m_1, m_2 \in \{1, 2, 3, 4\}$. Path-dependent hypothesis probabilities are evaluated for these two cases using (5), with the parent hypotheses $\bar{\lambda} \in \{\lambda_1^1, \lambda_1^2, \lambda_1^3, \lambda_1^4\}$ taken in turn, where $\bar{\tau}^{(m)} = \tau_1^{m_1}$ and with index $i = 2$. For independent Gaussian distributed radar track states, the track association likelihood is given by (8) with $\bar{y}^{(m)} = y_1^{(m_1)}$ and $i = 2$. Summing over propagation paths yields the two path-dependent hypotheses $\lambda_{11} = \{\tau_1 \cup \tau_2\}$ and $\lambda_{12} = \{\tau_1, \tau_2\}$ evaluated in ground coordinates, with respective probabilities

$$\Pr\{\lambda_{11}|Z_1^k \cup Z_2^k, D^k\} = \sum_{\substack{m_1, m_2=1 \\ m_2 > m_1}}^4 \Pr\{\lambda_{11}^{m_1 m_2}|Z_1^k \cup Z_2^k, D^k\}, \quad (17)$$

and

$$\Pr\{\lambda_{12}|Z_1^k \cup Z_2^k, D^k\} = \sum_{m_1, m_2=1}^4 \Pr\{\lambda_{12}^{m_1 m_2}|Z_1^k \cup Z_2^k, D^k\}. \quad (18)$$

Hypothesis, λ		Constraint	No. of hypotheses
Label	Targets		
$\lambda_{111}^{m_1 m_2 m_3}$	$\{\tau_1^{(m_1)} \cup \tau_2^{(m_2)} \cup \tau_3^{(m_3)}\}$	$m_1 < m_2 < m_3$	4
$\lambda_{112}^{m_1 m_2 m_3}$	$\{\tau_1^{(m_1)} \cup \tau_2^{(m_2)}, \tau_3^{(m_3)}\}$	$m_1 < m_2$	24
$\lambda_{121}^{m_1 m_2 m_3}$	$\{\tau_1^{(m_1)}, \tau_3^{(m_3)} \cup \tau_2^{(m_2)}\}$	$m_1 < m_3$	24
$\lambda_{122}^{m_1 m_2 m_3}$	$\{\tau_1^{(m_1)}, \tau_2^{(m_2)} \cup \tau_3^{(m_3)}\}$	$m_2 < m_3$	24
$\lambda_{123}^{m_1 m_2 m_3}$	$\{\tau_1^{(m_1)}, \tau_2^{(m_2)}, \tau_3^{(m_3)}\}$	—	64

Table 1. Path-dependent track-to-target association hypotheses for a cluster of three radar tracks and four propagation paths.

The target pdf for hypothesis λ_{11} is given by a Gaussian mixture, with each mixture component corresponding with a constituent path-dependent hypothesis, and with the component weights determined by the path-dependent hypothesis probabilities. That is,

$$p(y^{(m_1, m_2)}(k) | Z_1^k \cup Z_2^k, D^k) = \sum_{\substack{m_1, m_2=1 \\ m_2 > m_1}}^4 \Pr\{\lambda_{11}^{m_1 m_2} | Z_1^k \cup Z_2^k, D^k\} \mathcal{N}(y^{(m_1, m_2)}(k); \hat{y}^{(m_1, m_2)}(k|k), P^{(m_1 m_2)}(k|k)), \quad (19)$$

where $\hat{y}^{(m_1 m_2)}$ and $P^{(m_1 m_2)}$ are the fused state estimate and covariance given by (10) and (11) respectively, when the substitutions $\bar{y}^{(\bar{m})} \rightarrow y_1^{(m_1)}$ and $y_i^{(m_i)} \rightarrow y_2^{(m_2)}$ are made. These fused target pdfs are also illustrated in Figure 2.

When the third track τ_3 is considered, there are a total of 140 path-dependent track-to-target hypotheses which may be formed, namely the single target hypotheses $\lambda_{111}^{m_1 m_2 m_3}$, the hypotheses with two targets $\lambda_{112}^{m_1 m_2 m_3}$, $\lambda_{121}^{m_1 m_2 m_3}$ and $\lambda_{122}^{m_1 m_2 m_3}$, and the hypotheses with three targets, $\lambda_{123}^{m_1 m_2 m_3}$, with $m_1, m_2, m_3 \in \{1, 2, 3, 4\}$. These path-dependent hypotheses are tabulated in Table 1, together with the path constraints which apply in each case and the number of different hypotheses of each type which can be constructed. In applying the path constraint $m_i > \max(\bar{m})$, it is assumed that the variance in the ground state estimate is small compared with the differences between the estimates. In practice, this condition may not be satisfied, and the path assignment constraint must be relaxed to a constraint of the form $m_i > \max(\bar{m}) - \mu$ for some small integer μ , to retain additional plausible hypotheses. For example, propagation via paths with near-equal path length may result in multipath tracks having similar group range, so that strict marginal ordering may prohibit valid path-dependent hypotheses from being formulated.

Hypothesis probabilities follow from (5) with the parent hypotheses $\bar{\lambda}$ chosen from the set $\{\lambda_{11}^{m_1 m_2}, \lambda_{12}^{m_1 m_2}\}$ with $m_1 m_2 \in \{1, 2, 3, 4\}$. For the hypotheses $\lambda_{123}^{m_1 m_2 m_3}$ and $\lambda_{112}^{m_1 m_2 m_3}$, the transformed track $\tau_3^{(m_3)}$ is independent of the preceding tracks, and no fusion is required. The pdf for target $t_3 = 3$ is then given by the Gaussian mixture (4) with $j = 3$ and $M = 4$. For the hypotheses $\lambda_{111}^{m_1 m_2 m_3}$, $\tau_3^{(m_3)}$ is fused with $\bar{\tau}^{(\bar{m})} = \tau_1^{(m_1)} \cup \tau_2^{(m_2)} \in \lambda_{11}^{m_1 m_2}$. Likewise, for hypotheses $\lambda_{121}^{m_1 m_2 m_3}$ and $\lambda_{122}^{m_1 m_2 m_3}$, $\tau_3^{(m_3)}$ is fused with $\tau_1^{(m_1)} \in \lambda_{12}^{m_1 m_2}$ and $\tau_2^{(m_2)} \in \lambda_{12}^{m_1 m_2}$, respectively. Recursive hypothesis evaluation proceeds as in the binary track cluster case discussed above. In Figure 2, example path-independent hypothesis probabilities and target pdfs are illustrated for the three track cluster case. It is worth noting, for example, that improved target localisation arises from multipath track fusion, as shown by comparing the pdfs corresponding with the single target hypothesis sequence λ_1 , λ_{11} and λ_{111} along the top branch of the hypothesis tree.

4. MULTIPATH HYPOTHESIS EVALUATION: DYNAMIC CASE

The extension of hypothesis formation and evaluation to incorporate scene dynamics is of central importance in OTHR surveillance applications. Track cluster membership becomes dynamic, as conceivably associated tracks are initiated and terminated with the evolving scene. Target motion and changing ionospheric propagation conditions must be accounted for in the evaluation of association hypothesis probabilities and in the estimation of (possibly fused) target states.

Hypothesis trees can be used to construct track-to-target association hypotheses in the dynamic case.¹⁰ Fortunately, the number of path-dependent association hypotheses is only a function of the number of tracks in a cluster and the number of feasible propagation paths, and does not grow with time.¹¹ This constraint is imposed by the assumption that received detections associated with a given radar track propagate via the same path throughout the track lifetime. However, the initiation and termination of tracks in an existing track cluster introduces considerable complexity into the hypothesis management scheme. New tracks must be inserted into existing marginally ordered track clusters, while the termination of existing tracks may relax old path constraints. Dynamic hypothesis management is not discussed further here.

Consider the path-dependent track-to-target association hypotheses generated when the next target state $y^{(m_i)}(k+1)$ associated with ground track $\tau_i^{(m_i)}$ is taken at time $(k+1)$. If $\tau_i^{(m_i)}$ is a new track initiated at time $(k+1)$, then $\bar{Z}^k \cap Z_i^{k+1} = \emptyset$, and the new hypotheses arising from the new track are generated and evaluated as in the static case. The constraints on path assignment m_i follow from the marginal ordering rules for those hypotheses where the new track is to be fused with existing target states. If $y^{(m_i)}(k+1)$ is a continuation of a track $\tau_i^{(m_i)}$ already existing at time k , then the path assignment m_i is constrained to that assigned at all earlier times to the radar track τ_i . For those hypotheses where $\tau_i^{(m_i)}$ is to be fused with an existing (possibly fused) track $\bar{\tau}^{(\bar{m})}$, the generalisation of (6) to incorporate time evolution gives the fused state pdf

$$p(y(k+1)|\bar{Z}^k \cup Z_i^{k+1}, D^{k+1}, \tau^{(\bar{m}, m_i)}) = C^{-1} \frac{p(y(k)|\bar{Z}^k, D^k, \bar{\tau}^{(\bar{m})})p(y(k+1)|Z_i^{k+1}, D^{k+1}, \tau_i^{(m_i)})}{p(y(k+1)|\bar{Z}^k \cap Z_i^{k+1}, D^k, \bar{\tau}^{(\bar{m})}, \tau_i^{(m_i)})}, \quad (20)$$

where $\tau^{(\bar{m}, m_i)} = \bar{\tau}^{(\bar{m})} \cup \tau_i^{(m_i)}$. Since $\bar{Z}^k \cap Z_i^{k+1} = Z_i^k$, the denominator in (20) follows from the extrapolation

$$p(y(k+1)|Z_i^k, D^k, \bar{\tau}^{(\bar{m})}, \tau_i^{(m_i)}) = \int p(y(k+1)|y(k)) p(y(k)|Z_i^k, D^k, \bar{\tau}^{(\bar{m})}, \tau_i^{(m_i)}) dy(k), \quad (21)$$

where the state transition pdf $p(y(k+1)|y(k))$ is derived from the dynamic model used in the tracking filter, transformed to ground coordinates using (2). The substitution of Gaussian pdfs into (20), analogous to the derivation of (10) and (11) in the static case, yields the optimal fusion equations in the information form¹⁷

$$\begin{aligned} P^{(\bar{m}, m_i)}(k'|k')^{-1} &= P^{(\bar{m}, m_i)}(k'|k)^{-1} \\ &+ P^{(\bar{m})}(k'|k')^{-1} + P^{(m_i)}(k'|k')^{-1} - P^{(\bar{m})}(k'|k)^{-1} - P^{(m_i)}(k'|k)^{-1}, \end{aligned} \quad (22)$$

and

$$\begin{aligned} P^{(\bar{m}, m_i)}(k'|k')^{-1} \hat{y}^{(\bar{m}, m_i)}(k'|k') &= P^{(\bar{m}, m_i)}(k'|k)^{-1} \hat{y}^{(\bar{m}, m_i)}(k'|k) + P^{(\bar{m})}(k'|k')^{-1} \hat{y}^{(\bar{m})}(k'|k') \\ &+ P^{(m_i)}(k'|k')^{-1} \hat{y}^{(m_i)}(k'|k') - P^{(\bar{m})}(k'|k)^{-1} \hat{y}^{(\bar{m})}(k'|k) - P^{(m_i)}(k'|k)^{-1} \hat{y}^{(m_i)}(k'|k), \end{aligned} \quad (23)$$

with $k' = k+1$ for brevity. In (22) and (23), the predicted states and covariances follow from the track filter dynamic model, transformed to ground coordinates. The analogy here is with hierarchical multisensor track fusion, with synchronised update of the central fusion node, and where the fusion node has knowledge of the state dynamic model used by each sensor, so that the predicted states and covariances may be centrally constructed. Since multipath radar track states are estimated simultaneously, (22) and (23) yield the optimal fused estimate of the target ground state for the hypothesised propagation path set $\bar{m} \cup \{m_i\}$, even for the case of nondeterministic target motion when process noise may be significant.

The generalisation of the association hypothesis probabilities (5) to the dynamic case requires the evaluation of a track association likelihood which ideally accounts for track history. However, static hypothesis evaluation, which considers the transformed track states only at the most recent time, is useful for computational expedience in many circumstances. In Ref. 19, a generic track association likelihood was developed for hierarchical multisensor track fusion with nondeterministic target dynamics. In the notation of this paper, the likelihood that the ground track $\tau_i^{(m_i)}$ is associated with a (possibly fused) ground track $\bar{\tau}^{(\bar{m})}$ is given by

$$L(\tau^{(\bar{m}, m_i)}|\bar{Z}^k, Z_i^k, D^k) = \frac{l(\tau^{(\bar{m}, m_i)}|\bar{Z}^k \cup Z_i^k, D^k)}{l(\bar{\tau}^{(\bar{m})}|\bar{Z}^k, D^k) l(\tau_i^{(m_i)}|Z_i^k, D^k)}, \quad (24)$$




Figure 3. Schematic diagram of three radar tracks τ_j with $j \in \{1, 2, 3\}$ at three times $k = 1, 2, 3$, with track state estimates shown as a function of group range R . Example one-dimensional target pdfs are shown at each time step for the path-independent hypotheses, together with representative hypothesis probabilities. Plot scales and probability conditional arguments are omitted for clarity.

where $l(\tau|Z^k)$ denotes the likelihood of track τ described by the state sequence $\{y(1), y(2), \dots, y(k)\}$ estimated using the associated measurements $Z^k = \{z(1), z(2), \dots, z(k)\}$, and defined by the recursion

$$l_{n+1} = \begin{cases} \bar{\nu} \int p(z(1)|y(1))p(y(1)) dy(1) & n = 1, \\ l_n \int p(z(n+1)|y(n+1))p(y(n+1)|Z^n) dy(n+1) & n > 1, \end{cases} \quad (25)$$

where $\bar{\nu}$ is the *a priori* mean number of targets, and where it is assumed that the probability of detection is unity, so that there are no missed detections.

An example of dynamic hypothesis formation and evaluation is shown in Figure 3. A cluster of three radar tracks $\{\tau_1, \tau_2, \tau_3\}$ is shown, with τ_1 and τ_2 existing for all three sample times $k = 1, 2, 3$, while τ_3 is initiated at time $k = 2$ and joins the cluster. Track estimates $\hat{x}_j(k|k)$ with $j \in \{1, 2, 3\}$ are shown as filled squares, and the tracks are marginally ordered with increasing group range R . It is assumed here that the ionospheric model described in Appendix A is valid, so that four propagation paths exist. The time history of the tracks suggests that the most likely path-dependent association hypothesis is $\lambda_{112}^{m_1 m_2 m_3} = \{\tau_1^{(m_1)} \cup \tau_2^{(m_2)}, \tau_3^{(m_3)}\}$, for some choice of paths $\{m_1, m_2, m_3\}$ such that $m_2 > m_1$ and m_3 is unconstrained. This follows from the observation that τ_1 and τ_2 appear to have a similar positive group range rate, whereas τ_3 exhibits decreasing group range with time. Example one-dimensional ground state target pdfs are shown for the path-independent association hypotheses at each time step, together with representative hypothesis probabilities. Target motion is apparent from the pdf evolution illustrated; improved target localisation is shown to result from multipath track fusion within likely association hypotheses.

5. SUMMARY

A multihypothesis track fusion algorithm for multipath OTHR tracks has been outlined. The implementation of path-dependent track-to-target association hypothesis formation and evaluation has been examined with reference to a simple static example. Sources of track dependence have been identified and discussed. A preliminary treatment of scene dynamics in the multipath track fusion algorithm has been presented. Expressions for the optimal fused target pdfs have been given in the dynamic case, and a procedure for incorporating track history into the track association likelihood computation has been described. Further work is required to fully implement dynamic hypothesis evaluation, and to account for propagation path dependence.

REFERENCES

1. T. A. Croft, Sky-wave backscatter: a means for observing our environment at great distances, *Rev. Geophys.* 10, 73, 1972.
2. J. M. Headrick, Looking over the horizon, *IEEE Spectrum*, pp. 36-39, July 1990.
3. I. W. Dall and D. J. Kewley, Track association in the presence of multi-mode propagation, *Proc. International Conference on Radar 92*, IEEE Publication 365, pp. 70-73, 1992.
4. J. Zhu, R. E. Bogner, A. Bouzerdoum and M. L. Southcott, Application of neural networks to track association in over the horizon radar, *Proc. SPIE 2233*, pp. 224-235, 1994.
5. T. Kurien, D. Logan and W. P. Berry, Fusion of OTH radar data, *Proc. Seventh Joint Service Data Fusion Symposium*, pp. 221-235, 1994.
6. W. J. Yssel, W. C. Torrez and R. Lematta, Multiple relocatable over-the-horizon radar (ROTHR) track data fusion (MRTDF), *Proc. 9th Nat. Symp. Sensor Fusion*, pp. 37-45, 1996.
7. J. L. Krolik and R. H. Anderson, Maximum likelihood coordinate registration for over-the-horizon radar, *IEEE Transactions on Signal Processing*, 45(4): 945-959, April 1997.
8. R. H. Anderson and J. L. Krolik, Over-the-Horizon Radar Target Localization using a Hidden Markov Model Estimated from Ionosonde Data, accepted, Special OTH Radar issue of *Radio Science*.
9. G. W. Pulford and R. J. Evans, A multipath data association tracker for over-the-horizon radar, *IEEE Trans on Aerospace and Electronic Systems*, accepted, February 1998.
10. D. J. Percival and K. A. B. White, Multipath Track Fusion for Over-the-horizon Radar, *Proc. SPIE 3163*, pp. 363-374, 1997.
11. D. J. Percival and K. A. B. White, Multipath Coordinate Registration and Track Fusion for Over-the-horizon Radar, accepted, *Signal Proc. Special Issue for the DSTO/AFOSR Signal Processing Workshop*, Victor Harbor, South Australia, 1997.
12. N. S. Wheadon, J. C. Whitehouse, J. D. Milsom and R. N. Herring, Ionospheric modelling and target coordinate registration for HF sky-wave radars, *HF Radio Systems and Techniques*, IEE Conference Publication 392, pp. 258-266, July 1994.
13. E. Skafidas and D. J. Percival, Estimation of ionospheric state using a networked array of ionospheric sounders, *Proc. 4th International Symposium on Signal Processing Applications, Brisbane, Australia*, pp. 845-848, 1996.
14. E. Skafidas, D. J. Percival and R. J. Evans, Robust ionospheric density estimation using a vertical incidence sounder array, submitted to *IEEE Trans. on Aerospace and Electronic Systems*, May 1997.
15. C. Y. Chong, S. Mori, and K. C. Chang, 'Distributed multitarget multisensor tracking', in *Multitarget Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom, Ed., Norwood MA: Artech House, 1990, pp. 247-295.
16. M. E. Liggins, C. Y. Chong, I. Kadar, M. G. Alford, V. Vannicola and S. Thomopoulos, Distributed fusion architectures and algorithms for target tracking, *Proc. of the IEEE* 85, 1, pp. 95-107, 1997.
17. Y. Bar-Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, New York: YBS, 1995.
18. K. C. Chang, R. K. Saha and Y. Bar-Shalom, On optimal track-to-track fusion, *IEEE Trans. Aero. Electron. Sys.* 33, pp. 1271-1276, 1997.
19. S. Mori, K. A. Demetri, W. H. Barker and R. N. Lineback, A theoretical foundation of data fusion - generic track association metric, *Proc. Seventh Joint Service Data Fusion Symposium*, pp. 585-594, 1994.

APPENDIX A. Two-layer Stochastic Ionospheric Model

The most simple stochastic ionospheric model which admits multipath propagation is described here to illustrate the CR formalism outlined in Section 2.1. The model consists of two spherically-symmetric ionospheric reflecting layers, with altitudes described by the random variables h_1 and h_2 , respectively, with $h_1 < h_2$. Physically, Layer 1 may correspond to the ionospheric E-layer, with an altitude of $h_1 \approx 110$ km, while Layer 2 may describe the F-layer with $200 \text{ km} < h_2 < 550$ km. The ionospheric state at time k is then given by $w(k) = (h_1(k) \ h_2(k))'$. If the two layers are independent and Gaussian distributed, then $w(k) \sim \mathcal{N}(w(k); \hat{w}(k|k), P_w(k|k))$, with conditional mean $\hat{w}(k|k) = E\{w(k)|D^k\} = (\hat{h}_1 \ \hat{h}_2)'$ and covariance $P_w(k|k) = \text{cov}\{w(k)|D^k\} = \text{diag}(\sigma_1^2 \ \sigma_2^2)$, estimated from the sequence of ionospheric sounding measurements D^k . The covariances σ_1^2 and σ_2^2 describe the uncertainty in the height estimation for the two layers, which depend on the quality of the sounding measurements and fidelity of the model. It is assumed here that the ionospheric layers are independent, so that the cross-covariance components of P_w are zero. If Layers 1 and 2 correspond with E-layer and F-layer respectively, then representative uncertainties are $\sigma_1 = 5\text{km}$ and $\sigma_2 = 50\text{km}$.

If Layer 1 is partially reflecting with some fraction $0 \leq \alpha < 1$ of the transmitted signal energy penetrating to Layer 2 above, four propagation paths exist between a target and the radar sites, as tabulated in Table 2, with their corresponding path probabilities. These probabilities are simply derived from energy budget considerations, assuming isotropic transmission and signal reception, and assuming the target radar cross-section is independent of aspect angle.

For each propagation path m , an analytic expression for the CR transformation $T^{(m)}$ between a radar track state $x \in \mathcal{R}$ and its corresponding target ground state $y^{(m)} \in \mathcal{G}$ may be derived. In typical OTHR surveillance applications, state vectors with four components are used, so that $x = (R \ A \ \dot{R} \ \dot{A})'$ is transformed to $y = (r \ a \ \dot{r} \ \dot{a})'$, where R is the *group range* or *slant range*, A is the *apparent azimuth*, r is the *ground range*, a is the *true azimuth*, and where the over-dot denotes the time derivative. The components of x are estimated from radar detections by the tracking filter, with the group range R inferred from the signal time delay between transmission and reception, the apparent azimuth A estimated from the received signal direction of arrival, the group range rate \dot{R} estimated from the measured Doppler imparted to the received signal by the target motion, and the apparent azimuth rate \dot{A} estimated from changes in the received signal direction.

The estimated group range R is half the sum of the transmission signal group path and the received signal group path. With the simplifying assumptions that the transmitter and receiver are collocated, and where the effect of target height on the group paths may be neglected, it is straightforward to show that

$$R = [R_g(h_t, r) + R_g(h_r, r)]/2 \quad (26)$$

with

$$R_g(h, r) = 2 \left[R_e^2 + (R_e + h)^2 - 2R_e(R_e + h) \cos \left(\frac{r}{2R_e} \right) \right]^{1/2}, \quad (27)$$

where the model heights h_1 or h_2 are appropriately substituted for the height of the transmission signal propagation layer h_t and the height of the received signal propagation layer h_r for each propagation path $m = 1, 2, 3, 4$, and where with R_e the radius of the Earth. Assuming that ionospheric dynamics may be neglected so that $dh_t/dt = dh_r/dt = 0$, the time derivative of (26) yields the slant range rate

$$\dot{R} = \dot{r} \sin \left(\frac{r}{2R_e} \right) \left[\frac{R_e + h_t}{R_g(h_t, r)} + \frac{R_e + h_r}{R_g(h_r, r)} \right]. \quad (28)$$

Using simple trigonometry, transformation expressions for the apparent azimuth A and azimuth rate \dot{A} can also be derived. For completeness,

$$A = \sin^{-1} \left[\sin a \sin \left(\frac{r}{2R_e} \right) \frac{2(h_r + R_e)}{R_g(h_r, r)} \right], \quad (29)$$

and

$$\dot{A} = \frac{2(R_e + h_r)}{R_g(h_r, r) \cos A} \left[\dot{a} \cos a \sin \left(\frac{r}{2R_e} \right) + \frac{\dot{r} \sin a \cos(r/2R_e)}{R_e} - \frac{2\dot{r}(R_e + h_r) \sin a \sin^2(r/2R_e)}{R_g^2(h_r, r)} \right]. \quad (30)$$

Path m	Transmit layer	Receive layer	Path Probability $\beta^{(m)}$
1	1	1	$(1 - \alpha)^2$
2	1	2	$2\alpha(1 - \alpha)$
3	2	1	$2\alpha(1 - \alpha)$
4	2	2	α^2

Table 2. Path labels and path probabilities for the two-layer stochastic ionospheric model.

Equations (26) with (27) and (28)-(30) constitute the inverse CR transformation for each of the four propagation paths corresponding with the four possible substitutions of h_t and h_r with h_1 and h_2 . Analytic expressions for the direct CR transformation may be derived by inverting (26)-(30) for each path. For the $m = 1$ path which involves signal transmission and reception via Layer 1, inverting (26)-(30) with $h_t = h_r = h_1$ yields

$$y^{(1)} = \begin{pmatrix} r \\ a \\ \dot{r} \\ \dot{a} \end{pmatrix} = \begin{pmatrix} 2R_e \cos^{-1} \left[1 + \frac{h_1^2 - (R/2)^2}{2R_e(R_e + h_1)} \right] \\ \sin^{-1} \left[\frac{R \sin A}{2(R_e + h_1) \sin(r/2R_e)} \right] \\ \frac{R\dot{R}}{2(R_e + h_1) \sin(r/2R_e)} \\ \frac{R\dot{A} \cos A}{2(R_e + h_1) \cos a \sin(r/2R_e)} + \frac{\dot{r} \sin A}{\cos a} \left[\frac{1}{R} - \frac{R \cos(r/2R_e)}{4R_e(R_e + h_1) \sin^2(r/2R_e)} \right] \end{pmatrix} \quad (31)$$

as the implicit CR transformation. An explicit version of the transformation follows by eliminating r and a on the right side with the appropriate ground state component substitutions. The $m = 4$ path CR transformation follows from (31) by replacing h_1 with h_2 throughout; the mixed path cases $m = 2$ and $m = 3$ are not treated explicitly here, for the sake of brevity. The Jacobian matrices in (3) follow by taking the appropriate partial derivatives of the components of (31). Even for this most simple analytic model, the evaluation of the Jacobians is tedious, and the use of computer algebra software packages is recommended.

6.3 MULTIPLE TARGET LOCALIZATION AND TRACK ASSOCIATION FOR OVER-THE-HORIZON RADAR

Multiple Target Localization and Track Association for Over-the-Horizon Radar

Richard H. Anderson
Department of Electrical and Computer Engineering
Duke University, Durham, NC 27708-0291

Monday, February 23, 1998

1 Objective

This research project focuses on developing and evaluating a statistical approach to multiple target localization and track association for over-the-horizon radar. The goal is to provide improved localization accuracy and robust track association by explicitly accounting for the variability in down-range ionospheric parameters. The method employs a statistical model for the uncertainties in ionospheric propagation conditions to allow joint determination of the target ground locations, propagation paths, and track families from tracks of slant coordinate observations.

2 Motivation

Over-the-horizon (OTH) radars exploit the refractive and multipath nature of high-frequency (HF) ionospheric propagation to perform wide-area surveillance. Target detection, in the presence of noise and clutter, is accomplished by tracking returns in the slant coordinates, x , of delay or slant range, (d) , slant range rate or Doppler (\dot{d}) , horizontal wavenumber or slant azimuth, (b) and slant azimuth rate (\dot{b}) . Mode linking is the association of different slant tracks into target families based on their relative slant coordinates. Coordinate registration (CR) is the process of localizing each target by converting these slant coordinates to ground coordinates, n , of ground range (n) , ground azimuth (θ) , speed (v_s) and heading (v_h) [1, 2].

Target localization is conventionally achieved by generating a CR table comprised of propagation paths consistent with the observed slant coordinates using an ionospheric model estimated from a quasi-vertical incidence ionogram (QVI) and wide-sweep backscatter ionogram (WSBI) measurements. A simplified CR plot of slant range versus ground range obtained by raytracing through a multiple quasi-parabolic ionospheric model is shown in Figure 1. In some cases, a slant track observation uniquely corresponds to a single ground position (e.g. a slant range of 1900 km in Fig. 1). In general, however, a slant track observation can correspond to multiple ground locations that are associated with different propagation paths called raymodes (e.g. a slant range of 1400 km in Fig. 1). This mode ambiguity problem is further complicated in situations where multiple targets have similar slant coordinate measurements. In this case, two slant tracks can either correspond to two different targets or to a single target with slant tracks that correspond to two different raymodes. To resolve the mode ambiguity problem, mode linking is typically accomplished by a) performing coordinate registration using all possible propagation paths, b) comparing the ground coordinates to all existing target ground tracks and c) associating tracks with ground




Figure 1: Slant range vs. ground range for one ionospheric realization at 14 MHz

position, speed and heading variances below given thresholds, or where a track can belong to multiple families, associating it with the family with the minimum variance in ground coordinates [3, 4]. The key assumption in this conventional approach is that the down-range ionosphere is known very precisely, where in reality, only estimates of down-range ionospheric parameters are typically available. Not surprisingly, errors in the estimated down-range ionospheric parameters can seriously degrade target localization accuracy.

It should be noted that a probabilistic CR and track association method, called multipath track fusion and similar to the approach taken here, has recently been pursued by Percival and White [5, 6]. The method draws from multiple hypothesis approaches for multisensor tracking of several different targets.

The approach taken here is designed to improve target localization accuracy by statistically modeling uncertainties in the ionospheric propagation conditions. Joint maximum likelihood (ML) coordinate registration and track association (CRTA) consists of determining the most likely target ground coordinates, slant track families and mode associations over an ensemble of ionospheric conditions consistent with the data. For greater computational efficiency, the likelihood surface is approximated by a hidden Markov model (HMM) for the probability of a slant coordinate observation "sequence" given a hypothesized target location. The likelihood surface computed for the ML target localization is also incorporated into a multiple hypothesis method for determining the correct slant track associations.

3 Background

3.1 Single Target Localization Using a Hidden Markov Model

A statistical model for radar slant data can be obtained by mapping random ionospheric parameters through a HF radiowave propagation model. The conventional approach to the modeling of HF radiowave propagation through the ionosphere has been ray theory [7, 8]. For a given plasma frequency profile, raytracing may be used to determine the slant coordinates produced by a target at a particular ground position. When the ionospheric plasma frequency profile depends on uncertain random parameters, however, the slant coordinates for each ground position are themselves random variables. An example of this is shown in Figure 2 where the colors signify different raymodes. The objective of optimal coordinate registration and track association is to determine the target locations and track associations given the observed slant data and a model for their underlying




Figure 2: Slant range vs. Ground Range for 40 ionospheric realizations at 14 MHz

probability distribution. In addition to improving robustness to ionospheric uncertainty, statistical modeling also facilitates the use of less predictable observations, such as raymode amplitude rankings in the coordinate registration process. Although too unreliable for use with a deterministic ionospheric model, raymode amplitude rankings do contain important information which can be used by stochastic models to increase the probability of making correct mode associations.

Ordering the observed slant coordinates, for one target and one radar dwell, in terms of the signal-to-noise ratio (SNR) of each return, let \mathbf{x}_n denote the observation from the n th strongest return. Associating numbers with the different raymode types, let s_1 denote the mode number corresponding to the strongest return, s_2 the mode number of the second strongest return, and so on, until s_M is the mode number of the weakest return. The complete slant coordinate observation "sequence" is then given by

$$\mathbf{x}_n = \begin{bmatrix} d_{s_n}(\mathbf{r}) \\ \dot{d}_{s_n}(\mathbf{r}) \\ b_{s_n}(\mathbf{r}) \\ \dot{b}_{s_n}(\mathbf{r}) \end{bmatrix} + \epsilon \quad \text{for } n = 1, \dots, N \quad (1)$$

where ϵ represents the slant coordinate estimation error or "jitter" due to the smoothing tracker and is modeled by an independent zero-mean Gaussian distribution. Note that Equation (1) depends on both the sequence of raymode numbers, s_n , and on the mapping from ground location to slant coordinates, $[d_{s_n}(\mathbf{r}), \dot{d}_{s_n}(\mathbf{r}), b_{s_n}(\mathbf{r}), \dot{b}_{s_n}(\mathbf{r})]^T$, given the raymode number. Because both these quantities are random due to ionospheric variability, this statistical model is doubly stochastic. Given a specific set of ordered slant coordinates, $\mathbf{x}_1, \dots, \mathbf{x}_N$, the objective here is to determine the most likely target ground coordinates.

Let $p_{\mathbf{x},s}(\mathbf{X}, S|\mathbf{r})$ denote the joint probability density function of the ordered slant coordinate observations, $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]$, and corresponding raymodes, $S = [s_1, \dots, s_N]$, for a hypothesized ground location, \mathbf{r} . The maximum likelihood estimate (MLE) of ground location is obtained by substituting the observed slant coordinates into $p_{\mathbf{x},s}(\mathbf{X}, S|\mathbf{r})$ and maximizing with respect to raymode sequence, S , and hypothesized ground location, \mathbf{r} , i.e.

$$\tilde{\mathbf{r}}_{\text{ML}} = \arg \max_{\mathbf{r}} \left[\max_S (p_{\mathbf{x},s}(\mathbf{X}, S|\mathbf{r})) \right] \quad (2)$$

To reduce the computation associated with performing the maximization in Equation (2), a hidden Markov model (HMM) is used here to characterize \mathbf{X} [9]. In the terminology of HMM's, the

ordered slant coordinates define the *observation sequence* and their associated unobservable raymodes the *hidden state sequence*. The use of a HMM to compute the likelihood surface requires two approximations. First, given that the raymode sequence is known *a priori*, the slant returns of different raymodes are assumed to be statistically independent, i.e.

$$p_{\mathbf{x}|s}(\mathbf{X}|S, \mathbf{r}) = \prod_{n=1}^N p_{\mathbf{x}_n|s_n}(\mathbf{x}_n|s_n, \mathbf{r}) \quad (3)$$

This assumption is justified by the fact that the raymodes traverse different parts of the ionosphere exhibiting largely independent variations. Second, the raymode sequence is described by a first-order nonstationary Markov model such that:

$$P(S|\mathbf{r}) = \prod_{n=1}^N P_n(s_n|s_{n-1}, \mathbf{r}). \quad (4)$$

This implies that given the n th strongest raymode type, s_{n-1} and s_{n+1} are statistically independent. While clearly an approximation for $n > 2$, both simulation and real data results presented below suggest that this adequately models the structure of the raymode sequence. Substituting the product of (3) and (4) for $p_{\mathbf{x},s}(\mathbf{X}, S|\mathbf{r})$, the maximization in Equation 2 can be performed, via the Viterbi algorithm, by computing the set of path weights, $W_n(s_{n,k}, \mathbf{r})$ where $s_{n,k}$ denotes the k th value of s_n . The MLE of ground location is found by maximizing $W_M(s_{M,k})$ over k and a discrete grid of hypothesized ground coordinates. Note that, in contrast to the multipath track fusion approach, this iterative method ensures that the MLE grows linearly rather than exponentially with increasing N .

3.2 Multiple Target Mode Linking

The single target localization method can be extended to jointly perform target localization and track association in multiple target scenarios using individual radar dwells. Due to the potentially large number of slant returns and targets, the returns are first grouped into rough clusters based on their relative *slant* coordinates. In particular, the returns are grouped according to relative slant range rate, slant azimuth, and slant azimuth rate. The first cluster is centered about the coordinates of the strongest return and includes any returns within a specified window of \hat{a} , \hat{b} , and \hat{b} . The second cluster is centered about the next strongest return outside of the first cluster and the process continues until all returns have been assigned to a cluster. This clustering method also precludes associations of tracks with widely separated Doppler or slant azimuth observations.

Following the formation of track clusters, the track associations and target locations are jointly determined using multiple hypothesis testing for each cluster. Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_M]$ and $S = [s_1, \dots, s_M]$ now respectively denote the N SNR-ranked slant returns and corresponding raymodes within a given cluster. Let $\mathbf{R} = [\mathbf{r}_1, \dots, \mathbf{r}_L]$ denote the set of ground coordinates of L targets within the cluster. To determine the target family assignments, M hypotheses are formed using the leaf nodes of a hypothesis tree. The hypothesis tree is shown for $N = 4$ in Figure 3. Note that the different association hypotheses here correspond just to the leaf nodes and the tree structure is not dependent upon the particular multipath conditions. The raymodes' dependence, for a track family, is contained in the HMM transition and output probabilities and the most probable raymode sequence is given by the Viterbi algorithm.

To jointly determine $\tilde{\mathbf{R}}$, \tilde{H}_m , and \tilde{S} , maximize over

$$p(\mathbf{R}, S, H_m|\mathbf{X}) = \frac{p(\mathbf{X}, S|\mathbf{R}, H_m) p(\mathbf{R}) P(H_m)}{p(\mathbf{X})}, \quad m = 1, \dots, M. \quad (5)$$




Figure 3: Multiple Hypothesis Track Association Tree

The hypotheses are treated as equally likely, (i.e. $P(H_m) = 1/M$) and the joint *a priori* ground coordinate distribution, $p(\mathbf{R})$, is also modeled by a uniform distribution. The likelihood function, $p(\mathbf{X}, \mathbf{S} | \mathbf{R}, H_m)$, for the m th hypothesis can be rewritten as

$$p(\mathbf{X}, \mathbf{S} | \mathbf{R}, H_m) = p(\mathbf{X} | \mathbf{S}, \mathbf{R}, H_m) P(\mathbf{S} | \mathbf{R}, H_m) \quad (6)$$

To reduce the computational load associated with determining the likelihood function, the slant coordinate observations of different targets are assumed to be statistically independent, i.e.

$$p(\mathbf{X} | \mathbf{S}, \mathbf{R}, H_m) = \prod_{l=1}^L p(\mathbf{X}_l | S_l, \mathbf{r}_l, H_m) \quad (7)$$

where $\mathbf{X}_l = [\mathbf{x}_{1,l}, \dots, \mathbf{x}_{N,l}]$ and $S_l = [s_{1,l}, \dots, s_{M,l}]$ denote the slant coordinate observation and raymode "sequences", respectively, for the l th target and m th hypothesis. Further, the raymode assignments of different targets are also treated as statistically independent, i.e.

$$P(\mathbf{S} | \mathbf{R}, H_m) = \prod_{l=1}^L P(S_l | \mathbf{r}_l, H_m). \quad (8)$$

The joint *a posteriori* probability, $p(\mathbf{R}, \mathbf{S}, H_m | \mathbf{X})$, is then

$$p(\mathbf{R}, \mathbf{S}, H_m | \mathbf{X}) = \frac{\prod_{l=1}^L (p(\mathbf{X}_l | S_l, \mathbf{r}_l, H_m) P(S_l | \mathbf{r}_l, H_m) p(\mathbf{r}_l)) P(H_m)}{p(\mathbf{X})} \quad (9)$$

and the maximization can be performed term by term. To focus on minimizing large errors, the maximization is performed over the marginal probability $p(\mathbf{S}, H_m | \mathbf{X})$. The ML target position estimates,

$$\hat{\mathbf{R}}_{ML} = \max_{\mathbf{R}, \mathbf{S}} p(\mathbf{X}, \mathbf{S} | \mathbf{R}, H_m), \quad (10)$$

are then computed.

3.3 Estimating HMM parameters from Ionospheric Soundings

Estimating the parameters of a probabilistic model invariably requires a "training set" of measurements derived from data statistically similar to the observations of interest. In the current application, this means obtaining measurements of the ionosphere which encompass the variability

exhibited in the region roughly half-way between the radar and the target. The instruments typically available to obtain this data are the QVI and WSBI. The QVI is a measurement of group delay versus frequency for HF skywave propagation between a transmitter and receiver less than about 150 km apart. A typical daytime QVI measurement is shown in Figure 4, where the different traces discernable are the so-called ordinary and extraordinary rays resulting from the anisotropic polarization-dependent propagation through the ionosphere. Plots such as that of Figure 4 can




Figure 4: QVI (06Jun97-1806)

be inverted for the vertical electron density, or equivalently, plasma frequency profile near the radar. To simplify raytracing, the plasma frequency profile is then often represented by analytic layers (e.g. Chapman or quasi-parabolic E, F1, and F2) which are parameterized by their heights, critical frequencies, and thicknesses. To determine the down-range ionospheric layer parameters, the overhead profile parameters estimated from the QVI can be used as starting points for fitting the WSBI leading edge. The WSBI measures the ground backscatter intensity as a function of




Figure 5: WSBI (06Jun97-1806, Bearing = 171°)

time delay and frequency, as illustrated in Figure 5. The WSBI leading edge is an estimate of the minimum group delay of ground returns as a function of frequency [1]. Precise fitting of a set of WSBI leading edges at different azimuths requires raytracing through an ionospheric model whose layer parameters vary in range, altitude, and azimuth. However, because of the limited nature of the ionospheric data, estimates of the 3-D spatially-varying ionospheric parameters are not unique. Thus in addition to QVI and WSBI data, estimation of the down-range ionospheric model relies heavily on the use of empirical ionospheric models such as FIRIC, a combination of

the Fully Analytic Ionosphere Model and the International Reference Ionosphere using Chapman layers [11, 12]. The combination of limited sounder data and historical ionospheric modeling leaves considerable uncertainty in the estimated down-range profiles which motivates the use of the ML target localization method described above.

In light of the fact that computational or empirical models of ionospheric variability have not yet been available, the approach taken here is to treat the plasma frequency profile as a spatially homogeneous random process in latitude and longitude around the midpoint between the radar and the dwell illumination region (DIR). Adding the assumption of spatial ergodicity, samples of the 3-D ionospheric estimate in azimuth and ground range can be treated as different 1-D realizations of the midpoint profile. Let F denote the distribution of the the down-range profile parameter samples $\mathbf{g}(\mathbf{r}_i)$ taken at ground locations \mathbf{r}_i . In [13], F was assumed known, but in practice F must be estimated from the QVI and WSBI measurements. Smoothed bootstrap resampling, described below, is a means of generating realizations of \hat{F} without the need to explicitly estimate a complicated joint distribution function [14, 15]. Once a large set of random realizations from \hat{F} is available, the parameters of the HMM for each hypothesized target location can be determined by Monte Carlo evaluation of the raytrace propagation model. For example, an estimate of the discrete-valued transition probability, $P_{n|n-1}(s_n|s_{n-1}, \mathbf{r})$ can be computed by using the proportion of realizations such that mode s_n is the n th strongest mode given that s_{n-1} is the $(n-1)$ th strongest mode. Similarly, an approximation to the continuous-valued HMM output probability density function, $p_{\mathbf{x}_i|\mathbf{s}}(\mathbf{x}_i|s_n, \mathbf{r})$, is a Gaussian PDF with sample mean $\hat{\mu}_{s_n}(\mathbf{r})$ and covariance

$$\hat{\mathbf{C}}_{s_n}(\mathbf{r}) = \text{Cov} \left(\begin{bmatrix} d_{s_n}(\mathbf{r}, \gamma_i) \\ \dot{d}_{s_n}(\mathbf{r}, \gamma_i) \\ b_{s_n}(\mathbf{r}, \gamma_i) \\ \dot{b}_{s_n}(\mathbf{r}, \gamma_i) \end{bmatrix} \right) + \Sigma_\epsilon \quad (11)$$

where $[d_{s_n}(\mathbf{r}, \gamma_i), \dot{d}_{s_n}(\mathbf{r}, \gamma_i), b_{s_n}(\mathbf{r}, \gamma_i), \dot{b}_{s_n}(\mathbf{r}, \gamma_i)]^T$ are the slant coordinates for a target at ground location \mathbf{r} using the i th ionospheric realization. Note that in Equation (11), the first term depends on the ionospheric variability and the second term contains the slant coordinate jitter covariance, Σ_ϵ , determined by radar parameters such as SNR, coherent integration time (CIT) and bandwidth.

To determine the HMM parameters, realizations from \hat{F} must be determined from a limited number of 1-D samples of the midpoint profile parameters via the smoothed bootstrap resampling method. Smoothed bootstrap realizations of the parameter vectors can be generated by the following procedure:

1. Draw a bootstrap sample, $\mathbf{g}_1^*, \dots, \mathbf{g}_N^*$ of i.i.d. variates with distribution $F_{\mathbf{g}}$ by selecting N realizations *with replacement* from the original L sampled parameter vectors.
2. Generate N i.i.d. pseudorandom vectors $\mathbf{e}_1, \dots, \mathbf{e}_N$ with zero mean Gaussian distribution $F_{\mathbf{e}}$ and form $\Psi_k^* = (1 - \lambda)^{\frac{1}{2}} \mathbf{g}_k^* + \lambda^{\frac{1}{2}} \mathbf{e}_k$ where $0 \leq \lambda \leq 1$.

These steps can be repeated to generate a large number of ionospheric realizations. In order to ensure that the bootstrap realizations have the same second-order statistics as the sampled parameter vectors, the covariance matrix of $F_{\mathbf{e}}$ is set to $\mathbf{C}_{\mathbf{e}} = \hat{\mathbf{C}}_{\mathbf{g}}$.

4 Preliminary Results

4.1 Single Target Scenario

The random ionospheric model in the single target simulation and real data results was based on QVI and WSBI measurements, taken December 7 and 8, 1994. First, overhead and down-range ionospheric parameters of a 3-D ionospheric model were determined using software routines developed by Nickisch and Hausman [11]. The 3-D model down-range profile parameters were sampled at eight ranges along the eight bearings to generate 64 ionospheric samples over roughly an 800x800 km down-range sample region represented by the grid crossings in Figure 6. To reduce




Figure 6: Sampling region of random ionospheric model

the computational burden of Monte Carlo evaluation of the raytracing model, an analytic multi-quasi-parabolic (MQP) raytracing model was employed. For the results here, the MQP layer heights and critical frequencies were taken equal to the 3-D model's parameters, and the semi-thickness parameters were computed from the layer heights to approximate the vertical profiles of the 3-D model. After obtaining samples from the 3-D model, the quasi 2-D MQP zenith tilt angle parameters were computed for each sample by a least-squares fit of the MQP synthetic leading edges to the corresponding WSBI leading edges. Comparisons of CR' curves suggest that the random realizations of the quasi 2-D MQP statistical model both encompass and closely approximate the results of numerical raytracing through the 3-D ionospheric model.

200 ionospheric realizations, generated from the down-range samples via smoothed bootstrap resampling with $\lambda = 0.01$, were used in Monte Carlo evaluation of the MQP raytracing propagation model to compute the slant ranges and associated raymode amplitudes for the bistatic extraordinary raymodes. The random propagation model was run at 22 MHz for a scenario with a minimal number of possible paths between the radar and target and also was run at 12 MHz for a case with a large number of raypaths. IIMM parameters for the 12 strongest raymodes were retained for ground ranges between 2000 and 2500 km in 1 km intervals. The *a priori* statistical information contained in the raymode amplitudes can be appreciated by considering the probability that a particular raymode corresponds to the strongest return. The probabilities for the raymode types, are plotted as a function of ground range in Figure 7. Note that at 22 MHz the probabilities have little variation in ground range. In contrast, the probabilities for 12 MHz indicate that more raymodes are probable for the strongest observation. The changes in these probabilities correspond to the ground ranges over which the F1 layer caustic exists in the transmit or receive paths.

A conventional CR method used for comparing localization accuracy was defined by the average of the ground ranges corresponding to raymodes identified as having the minimum ground range variance. This method is referred to here as the minimum variance (MV) CR method. The ionospheric model assumed by the conventional model was that obtained using the mean MQP




Figure 7: Probability of Different Raymode Types for Strongest Return

parameters from the ionosonde data. For 50 Monte Carlo trials, ground range estimation was performed for a target at each ground range in 10 km intervals between 2000 and 2500 km using a maximum of 3 observations. A zero-mean Gaussian slant range jitter component, with $\sigma_\epsilon = 3.0$ km, was added to model delay estimation error. A comparison of the MLE performance to the conventional MVE method is the variation in average absolute miss distance versus ground range, shown in Figure 8. Observe that the average miss distance achieved by the ML method




Figure 8: Average Absolute Miss Distances

is consistently less than that of the MV method. Also note that, in contrast to the MV, the ML method exploits the additional raymodes at 12 MHz to achieve better performance than for a single mode at 22 MHz. The performance improvement of the ML method over the MV method varies with true ground range but can be as much as 5 times more accurate.

To determine the achievable ML CR accuracy for a single target on real data, the method was tested on radar returns from a beacon at a ground range of 2192.66 km. The likelihood function, on a grid of hypothesized ground ranges, is shown in Figure 9 for 5 beacon dwells. Note




Figure 9: Likelihood function for 5 beacon dwells

that each likelihood function realization has multiple peaks indicating that there were several ambiguities in the ground range estimates. The ambiguities arise from the raymode uncertainty and the secondary peaks in the likelihood function correspond to less probable raymodes for a single return. To compare the ML method with conventional methods, histograms of the range errors, shown in Figure 10, were computed for each method using over 250 dwells recorded during 14 different tracking periods. The miss distances are reported in normalized units. The top two




Figure 10: Beacon Ground Range Error Histograms for Dec 7-8, 1994 with over 270 dwells, 8 slant range outliers removed and a true ground range of 2192.66 km

histograms correspond to the conventional methods of the existing radar system and the 3-D model respectively. The average absolute miss distances (AVMD) and average ground range biases for the three methods are shown at the left sides of the histograms. The ratios of the average absolute miss distances indicate that the ML method offers nearly a 2 to 1 improvement over the deterministic 3-D model and nearly a 3 to 1 improvement over the CR method in the existing radar system.

4.2 Multiple Target Scenario

To illustrate the potential performance improvement of the proposed CRTA method over conventional CR and mode linking approaches in multiple target scenarios, localization results are reported here for ground-truthed real data collected on June 6, 1997. The data, collected through a collaborative effort between Duke and Naval Research Laboratory, was part of an experiment to localize and track a Piper Aztec aircraft as well as commercial aircraft flying between Florida and Puerto Rico. Each method's performance was evaluated using both simulated OTH radar data, synthesized from FAIR ground-truth based on line-of-sight radar data, and real OTH radar slant track data of commercial flight RICH4814, shown in Figure 13.

The 3-D ionospheric model was determined as described in section 4.1. However, the statistics of the down-range ionospheric parameters were computed in a slightly modified approach. The ionospheric parameter correlation coefficients were computed from a six hour time-series of QVIs. The corresponding covariance, C_g , was determined from the correlation coefficients and parameter variance estimates from five most recent overhead profiles. Smoothed bootstrap resampling, with $\lambda = 0.1$ using samples from the 2 most recent WSBLs, was applied to generate 200 down-range profile realizations. HMM parameters were computed for the 9 most probable bistatic extraordinary raymodes via Monte Carlo evaluation of the MQP raytracing propagation model.

For 100 Monte Carlo trials, the coordinate registration and track association were performed using a maximum of 2 synthesized slant tracks of 20 ground-truth positions of flight RCH4814. The histograms of target ground range errors are plotted as a function of true ground range in Figure 11. Note that the number of estimates at a given range bin is indicated by the scale on the

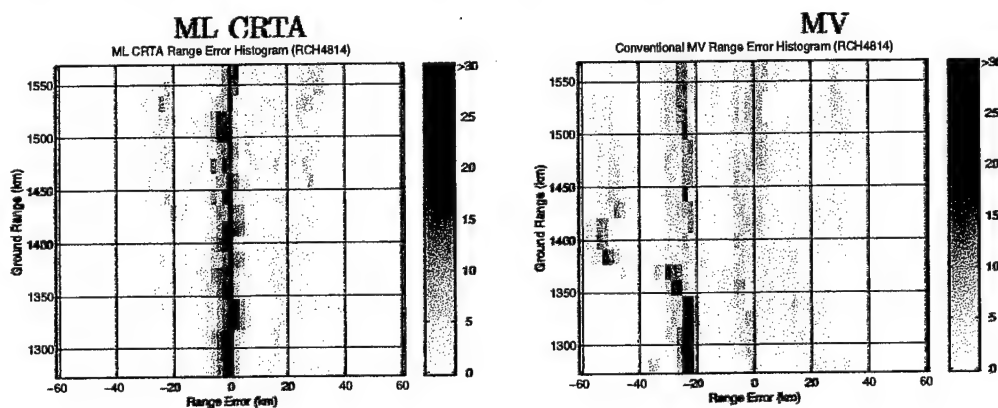


Figure 11: Ground Range Error Histograms

right of each histogram. An ideal histogram of ground range error would be a straight vertical line at 0 km. Note that the MV histograms are widely spread for many ground ranges and exhibit a consistent multimodal characteristic that results from erroneous mode associations. In contrast, the ML CRTA histograms are more tightly grouped around 0 km range error and has much smaller secondary modes. The average absolute miss distances, shown in Figure 12, further indicate the performance of the two methods. The performance improvement varies with ground range, but




Figure 12: Average Absolute Miss Distances for Simulation of RCH4814

the ML CRTA method can be as much as 3 to 5 times more accurate.

The CRTA method has been tested on real slant track data between 1806 and 1830 GMT, on June 6, 1997. The FAA ground-truth indicates that the target density for this DIR during this time was moderate and also that many of the flights had headings tangential to the propagation paths resulting in Doppler frequencies that were masked by ground clutter returns. The results presented here correspond to an aircraft along a fairly radial heading with Doppler well separated from ground clutter returns. The CRTA method was applied to all the slant tracks generated by the radar for this DIR and time period. Both the existing mode linker and the CRTA method associated three separate slant tracks (28588, 28596, 28613) with a single target ground track (40327). The ground

position estimates of the two methods and the FAA' ground-truth are overlaid in Figure 13. The




Figure 13: Ground Tracks of flight RCH4814

enlarged symbols and labels at the upper left indicate the final ground positions for these times. Note that the radar's estimates were initially more accurate, but the CRTA track soon approached the ground-truth positions. The radar's initial position estimates are the "fused" results from 2 OTHR radars with overlapping coverage of flight RCH4814, whereas the CRTA method is only applied to data from one OTHR radar. The absolute miss distances for the two methods were computed for each revisit and are shown in Figure 14. These results suggest the proposed ML CRTA method can potentially provide improved target localization accuracy.




Figure 14: Average Absolute Miss Distances for flight RCH4814

5 Future Work

To date, the performance of OTH radar target localization and track association methods has not been evaluated on a common and extensive set of radar data. Through continued cooperation with the Naval Research Laboratory, a "testbed" data set is proposed to be obtained for use in determining the performance of target localization and track association methods including the method described here. Ideally, this data set would include real radar data and ground-truth data collected for many different radar waveforms, ionospheric conditions, target densities, times, dates, and geographic locations. The anticipated radar data types include amplitude-range-Doppler maps, peak detections, slant tracks of existing tracking filters, and radar ionosonde data (i.e. QVIs and

WSBIs). The ground-truth data may consist of a combination of line-of-sight radar data and global-positioning-system (GPS) data. We propose to augment the real tracking data with simulation data based on ground-truth position reports and radar ionospheric soundings. Further, formation of a set of standard performance metrics (e.g. miss distances, error histograms, probabilities of correct localization and probabilities of correct association) is proposed to evaluate the different methods.

Several expansions of the current method are proposed to achieve greater robustness to ionospheric uncertainties and improvements in target localization accuracy. First, the approach may be extended to perform block processing of slant tracks over multiple revisits. One possibility involves incorporating a model of target dynamics by treating the ground positions as the states of a Markov chain and the likelihood functions as output probability densities. A Viterbi tracking algorithm could then be applied to a time-block of several revisits and the track association hypotheses evaluated for the whole block instead of individual radar dwells. The formation of a dynamic hypothesis tree structure is an alternative means of incorporating previous track association information into the current hypothesis evaluations. Second, an extension for rejecting or suppressing false tracks is proposed. The track association hypothesis tree could be modified to include a null hypothesis for each slant track to indicate when it corresponds to only a series of peaks due to noise or clutter. Third, we propose to extend the method to perform track association for asynchronous slant tracks generated by different OTH radars with overlapping dwell illumination regions. If the radars' dwell times were synchronized, the joint likelihood surface can be treated as a separable function of independent likelihood surfaces computed from different OTH radars' data. This principle will be combined with the multiple hypothesis testing and block Viterbi tracking approaches to perform track association of tracks originating from multiple OTH radars.

References

- [1] L. F. McNamara, *The Ionosphere: Communications, Surveillance, and Direction Finding*. Krieger Publishing, 1991.
- [2] N. S. Wheldon, J. C. Whitehouse, J. D. Milsom, and R. N. Herring, "Ionospheric modelling and target coordinate registration for HF sky-wave radars," in *IEE Sixth International Conference on HF Radio Systems and Techniques (IEE Conf. Publ. No. 392)*, (York, UK), pp. 258-266, 1994.
- [3] E. J. Ferraro, *ROTHR Mode Linking*. Raytheon Electronic Systems, August 1995. Memo Number EJJ:95:13.
- [4] Raytheon Company, *Relocatable Over The Horizon Radar - System Design and Operations Document (Parts I and II)*, June 1992. Report prepared under contract N00039-90-C-0027.
- [5] D. J. Percival and K. A. B. White, "Multipath track fusion for over-the-horizon radar," *Proc. SPIE 3163*, September 1997. To appear.
- [6] D. J. Percival and K. A. B. White, "Multipath coordinate registration and track fusion for over-the-horizon radar." Submitted to Elsevier Preprint, September 1997.
- [7] R. M. Jones and J. J. Stephenson, "A versatile three-dimensional ray tracing computer program for radio waves in the ionosphere," tech. rep., Office of Telecommunications, Boulder, CO, October 1975.

- [8] K. Davies, *Ionospheric Radio*. London, UK: Peter Peregrinus Ltd., 1990.
- [9] A. B. Poritz, "Hidden markov models: a guided tour," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 1, pp. 7-13, April 1988.
- [10] H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I*. New York: John Wiley and Sons, 1968.
- [11] L. J. Nickisch and M. A. Hausman, "CREDO: Coordinatic registration onhancement by dynamic optimization," in *Ionospheric Effects Symposium*, (Alexandria, VA), pp. 1B-2-1 - 1B-2-9, May 1996.
- [12] C. J. Coleman, "A propagation model for HF radiowave systems," in *MILCOM*, pp. 875-879, 1994.
- [13] J. L. Krolik and R. H. Anderson, "Maximum likelihood coordinate registration for over-the-horizon radar," *IEEE transactions on Signal Processing*, vol. 45, pp. 945-959, April 1997.
- [14] B. Efron, *The Jackknife, the Bootstrap and other Resampling Plans*. Philadelphia, PA: SIAM, 1982.
- [15] R. H. Anderson and J. L. Krolik, "Over-the-horizon radar target localization using a hidden markov model estimated from ionosonde data," *Radio Science*, December 1997. To appear.

March 1998 Australia Trip Report

Richard H. Anderson
Department of Electrical and Computer Engineering
Duke University, Durham, NC 27708-0291
Advanced Radar Systems, Naval Research Laboratory
Washington DC 20375

April 13, 1998

This report is intended to document work by Richard Anderson from March 9-19, 1998 associated with March 1998 US/AS MOA meeting on OTH radar. Participation in the coordinate registration experiment at IRSU in Alice Springs, technical discussions of multipath track association/track fusion approaches, and both the current status and future work on target localization and track association are all summarized here.

1 Objectives

To further the research on OTH radar target localization and track association, three objectives were identified for the March 1998 MOA meeting. First, Matlab m-files for both the Maximum Likelihood Coordinate Registration and Track Association (ML CRTA) method and Matched Field Altitude Estimation (MFAE) algorithm, developed by Mike Papazoglou and Jeff Krolik, were to be shared with Wide Area Surveillance Division (WASD) personnel at the Defence Science and Technology Organisation (DSTO) in Salisbury, SA. Second, participation in the Coordinate Registration (CR) data collection and evaluation during the week of March 9 was arranged. Third, plans were made for technical discussions of ML CRTA and Multipath Track Fusion (MPTF) with John Percival and Kruger White.

2 Work Accomplished

- **Matlab Code:** Both the ML CRTA and MFAE algorithms in Matlab were transferred to a Pentium personal computer, provided by DSTO, on March 9 at IRSU. The code was placed in the C:\Duke directory. Both algorithms successfully ran to completion on ROTH data examples in initial testing. Associated executables of Mission Research CREDO software routines were successfully installed by L. J. Nickisch in the C:\credo directory. It was anticipated that the PC would be returned to Salisbury during the week of March 23.
- **CR Experiment:** CR data was collected from March 9-13 for regions near the Darwin beacon using both the JFAS full array and half array configurations. Darwin microwave track data was also collected for use in WASD data fusion research and use as ground-truth data. Longer range CR data for several inbound Qantas flights from Bangkok to Sydney was also gathered but will remain classified. The work on ML CRTA in Alice Springs, following

the transfer of the Matlab code, consisted of developing/testing code to read JFAS slant track ASCII data files. It should be noted that significant effort was required of Rod Barnes, Thang Hoang and Peter Den Hartog to convert the ionogram and slant-track ASCII files from JFAS binary data. Contemporaneous ionogram and slant track data was collected on March 10 for just beacon tracks and on March 12-13 for all the slant tracks. On the 13th, the ML CRTA algorithm was applied to March 12 slant tracks from tasks used to track one of the Qantas flights. Time did not allow for comparison of the resulting ground tracks and ground truth to determine the method's performance. Since the PC provided by DSTO at 1RSU was unavailable in Salisbury from March 16-19, time was spent on a secure workstation in WASD to identify sets of slant tracks in interesting mode-linking scenarios. The table below indicates a small set of slant tracks collected on March 12th that has associated Darwin microwave ground-truth positions and potentially could be declassified and released to the US.

JFAS slant tracks near Darwin on March 12, 1998

Region/Task	Time (UT)	Track IDs
A	0444-0451	2749 2769 2841 2849 2854 2856 2866 2883 2898 2935 2936 2937 2938 2941 2963 2969 2970
B	0444-051	2889 2906 2917 2921 2927 2929 2931 2943 2944 2948
A	0501-0521	2856 2866 2929 2974 2977 3000 3007 3017 3066 3126 3129 3130 3131 3132 3149 3158 3172 3185 3187 3192 3199 3217 3225 3759
B	0501-0521	2889 2929 3024 3027 3149 3158 3159 3246 3255

This slant track data along with the contemporaneous ionograms, Darwin microwave tracks and records of the JFAS tasks' radar parameters (e.g. frequency, bandwidth, WRF, CIT, and full-array or half-array mode) would be useful for additional unclassified CR and track association research.

- **Technical Discussions:** Several discussions of the details of ML CRTA and MPTF were conducted with John Percival, Kruger White, and Branko Ristic. First, a Matlab implementation of MPTF was demonstrated by Kruger White. Second, some of the expressions and concepts of MPTF were clarified. These included the recursive hypothesis evaluation of the decision tree using a binary fusion term and also the computation of normalization constants. John and Kruger also identified the specific mode-linking/track-fusion problems that they are addressing as well as separate CR issues being addressed by Rod Barnes. It was noted that the current implementation of MPTF uses the static hypothesis evaluation and uses CR advise where either the raymodes are treated as equally probable or the raymode probabilities are proportional to the raymode amplitudes.

The remainder of the discussions centered on the ML CRTA algorithm and several significant issues were raised. Rough comparison of elapsed time for MPTF and ML CRTA indicated the larger computational load associated with the latter. The raymode amplitude ranking in ML CRTA may not distinguish different raymodes when significant overlap in the amplitudes'

Rician probability density functions exists. Different slant tracks from a single target can have significant correlations due to common process noise, common reflection heights on transmit or receive paths and shared peak detections. These dependences are not exploited in the current ML CRTA method. Further, the current ML CRTA method processes each revisit separately and the raymode assignments are not explicitly constrained to remain constant from one revisit to the next. This can result in a discontinuous ground track when a strong slant track starts later than a weak slant track assigned to the same target.

As an aside, note that on Wednesday March 18, the MFAE method was demonstrated on ROTH data (November 4th Aztec flight) for Justin Praschifka, Mike Turley and Branko Ristic and implementation of MFAE on JFAS data was briefly discussed.

3 Current Status

Presently, the existing ML CRTA method has been applied to real ROTH and JFAS data in multiple target scenarios. Preliminary performance has been reported for ROTH data corresponding to 1 FAA ground track from June 6, 1997. The Matlab implementation of ML CRTA as well as the Matched Field Altitude Estimation demo have been delivered to WASD at DSTO. A small set of JFAS radar data, collected during the March 9-13 CR experiment, is described here and suggested for possible declassification and release to the US. John Percival plans to visit Duke University on April 21 and to provide a copy of the MPTF Matlab code at that time.

4 Future Work

Several modifications and extensions of the current ML CRTA method are proposed to achieve improved target localization and track association performance. The first addition or modification of the ML CRTA method is block processing of slant track data to incorporate the slant tracks' time dependence and to constrain mode assignments over multiple revisits. One possible approach consists of using peak detections and performing target tracking in ground coordinates. Second, we propose to investigate the correlations of different slant tracks and targets and potentially modify the ML CRTA method to exploit these dependences. A third possible extension of the method is data fusion of tracks from overlapping DIRs and tracks from different OTH radar sites. We expect to evaluate these modifications and enhancements both on simulated data derived from real multiple target scenarios and on real ROTH and declassified JFAS radar data, if possible.

6.4 SECTION REFERENCES

- Yssel, W. and W. Torrez. 1998. "Recent Developments in Fusing Microwave Radar Tracks with Relocatable Over-the-Horizon Radar (ROTHR) Tracks," *Proc. Internat. Conf. Multisource-Multisensor Inform. Fusion (FUSION 98)*, pp. 757-764, Las Vegas, NV, (July).
- Percival, D. and K. White. 1997. "Multipath Track Fusion for Over-the-Horizon Radar," *Proc. SPIE Conf. Signal and Data Proc. of Small Targets*, pp. 363-374, San Diego, CA, (July).
- Percival, D. and K. White. 1997. "Multipath Coordinate Registration and Track Fusion for Over-the-Horizon Radar," *Proc. DSTO/AFOSR Signal Proc. Workshop*, Victor Harbor, AS, (June).
- Percival, D. and K. White. 1998. "Multihypothesis Fusion of Multipath Over-the-Horizon Radar," *Proc. SPIE Conf. Signal and Data Proc. of Small Targets*, Orlando, FL, (April).
- Anderson, R. 1998. "March 1998 Australia Trip Report," Duke University, Durham, NC, (April).
- Anderson, R. 1998. "Multiple Target Localization and Track Association for Over-the-Horizon Radar," Duke University, Durham, NC, (February).

7. CONCLUSIONS AND A VISION OF THE FUTURE

7.1 GOODMAN-MAHLER APPROACH TO GLOBAL DATA FUSION

Data Fusion Algorithms (such as the one that will likely fuse ROTH tracks with the Multiple Input Tracking and Control Subsystem [MTRACS] surveillance system tracks described in section 6.1) will rely increasingly on automatic, rule-based expert systems that incorporate both kinematic and attribute data as described previously. R. Mahler (references 1 and 2) has described a method, known as global data fusion, that is based on I. R. Goodman's random set theory (*cf.* reference 3), which generalizes the classical statistical approaches, and allows the combining of prior information of attributes with geodynamical data. This method is consistent with Bayesian and Neyman-Pearson statistics and makes no Gaussian or independence assumptions. Furthermore, this approach encompasses "ambiguous" classification evidence (such as fuzzy, Dempster-Shafer, rule-based, or partially probabilistic).

7.2 MATHEMATICAL FORMULATION OF GLOBAL DATA FUSION

What we seek is a way to represent a collection of targets as a single "global" target and a suite of sensors (similar or disparate) as a single "global" sensor. In this way, we can directly apply the standard results of statistical estimation theory for quite general data fusion problems. This is done by representing the collective statistical behavior of the sensors as a single "global" probability density function (PDF). In order to do this, one quickly realizes that the usual mathematical model for applied data fusion—the random variable (or more generally, the random vector)—is no longer adequate. In the Goodman-Mahler approach, the random vector is replaced by a novel kind of random variable called the random set, as pioneered in data fusion applications by I. R. Goodman of SSC San Diego. Using this notion, (as well as the Radon-Nikodym derivative to obtain the notion of a "global" PDF as proposed by R. Mahler), it turns out that virtually all of classical statistics can be extended to the multisensor, multitarget realm. Surprisingly, this is accomplished even if it is assumed that our sensors supply classification as well as positional data about targets. Global data fusion is based on so-called global PDFs. Global PDFs generalize ordinary PDFs to the case when the object to be estimated is a set $\otimes^* = \{\otimes_1, \dots, \otimes_t\}$ of targets being observed by a suite $\Theta^* = \{\Theta_1, \dots, \Theta_s\}$ of (not necessarily independent) sensors of varying kinds. It is assumed that the detectability of targets by sensors can vary with parameters such as range, aspect angle, and identity of targets. As it was stated above, the basic theoretical innovation that underlies the global data fusion is random set theory. A finite random subset of state space is simply a random variable on the set of all finite subsets of state space. In the same way that the statistics of an individual sensor can be modeled by a random vector, the statistics of the global sensor Θ^* can be modeled by a finite random set Σ whose values are "clustered" in some specified parameterized sense around the global truth state \otimes^* . Furthermore, just as the statistical behavior of a random vector can be specified in general by a probability measure, so the statistics of Σ can be specified by the "belief measure" defined by $\beta_\Sigma(S) = p(\Sigma \subseteq S)$. The value $f_X(x)$ of a conventional random vector X is a measure of how likely it is that the vector x occurs as a value of X . Likewise, the value $f_\Sigma(T)$ is a measure of how likely it is that the finite set T of position vectors occurs as a value of the random set Σ . Given this, the reasoning used in classical statistical estimation theory can be applied to multitarget, multisensor data fusion problems. That is, such algorithms can be mathematically modeled as black boxes, which sample the output of Θ^* and arrive at an estimate $J(\Sigma_1, \dots, \Sigma_m)$ of "global" ground truth \otimes^* . They accept independent "samples" S_1, \dots, S_m of corrupted sensor data as input. These samples are sets of vectors with varying numbers of elements. The global PDF f_Σ is assumed to have the form $f_{\Sigma|\theta}(s|\theta)$ where $\theta = (\theta_1, \dots, \theta_p) \in (R^n)^p$.

is a vector parameter that defines the mathematical form of f_{Σ} . Then a "global estimator" is a random vector of the form $J = J(\Sigma_1, \dots, \Sigma_m) \in (R^n)^p$ where $\Sigma_1, \dots, \Sigma_m$ are statistically independent random sets that have f_{Σ} as their global PDFs. Now by introducing discrete state variables to represent targets or target classes, one can extend this reasoning to data fusion problems that involve classification as well as tracking evidence. A discrete state variable is incorporated into the system state vector, which takes its values from the set of characterized targets. States of the system therefore have the general form (z, u) where z is a geopositional vector and u is an element from the set of characterized targets.

As an application of this approach, the two independent (but, under current requirements, soon to be interoperable) systems of particular interest for this research are the ROTH system and MTRACS. The current MTRACS display provides the operator with a surveillance picture based only on the fused kinematic aspects of the target measured by a net of microwave radars. In the final implementation of this system, ROTH tracks formed in regions of interest (but likely outside the range of the microwave radars) will be fused with the microwave radar tracks. Since it is likely that noncommensurate error ellipses will make straightforward track association between the ROTH tracks and the microwave radar tracks of the MTRACS system impractical, it will be necessary to design DFAs that exploit both kinematic (e.g., tracklets, origin of sensors, times of intercepts, slant or ground coordinates, range-rates, etc.) and attribute information (e.g., sensor assigned track numbers, target type identification, image, size, shape, etc.). In effect, the attribute cues could, if properly fused with the kinematic ones, provide a better Level 1 fusion product (situation picture or object refinement) for the MTRACS operator, as well as possibly enhance the quality of Level 3 fusion (threat assessment or threat refinement). The global data fusion approach provides a mathematical model for the quantitative evaluation of such data fusion algorithms. By using this approach to define the global statistics of such DFAs, it will be possible to evaluate measures of effectiveness for target detection and tracking, thus allowing designers to optimize the performance characteristics of candidate DFAs and to assess the merits of competing DFAs.

7.3 SUMMARY AND CONCLUSIONS

The methods described in this paper will enable a mathematical evaluation of the performance characteristics of quite general DFAs, that is, those that combine both classification and positional information in a multitarget and multisensor environment. This approach introduces the methodology to blend diverse measurements to form "best" estimates. Therefore, this work is oriented to the applied rather than the theoretical aspects of optimal estimation of global data fusion. In particular, this work has application in the case of fusing microwave radar tracks output by the MTRACS system with tracks of targets of interest generated by the ROTH system—both systems having tremendous capabilities as independent sources. Their interoperability and fusion will give the user a surveillance capability of tremendous coverage and flexibility, if properly designed. Much is known about the mathematical properties of optimal design in the classical case of a single sensor and a single target (*cf.* figure 7.3.1 and references 4–6). Similarly, in the case of multiple, nonhomogeneous sensors and a single target, many partial results are known. However, in today's information rich, highly connected computing environments, there still has been little, if any, investigation of the performance characteristics of combining information derived from a multitarget and multiple nonhomogeneous sensors. The approach described in this paper will play a key role in the design of optimal DFAs.

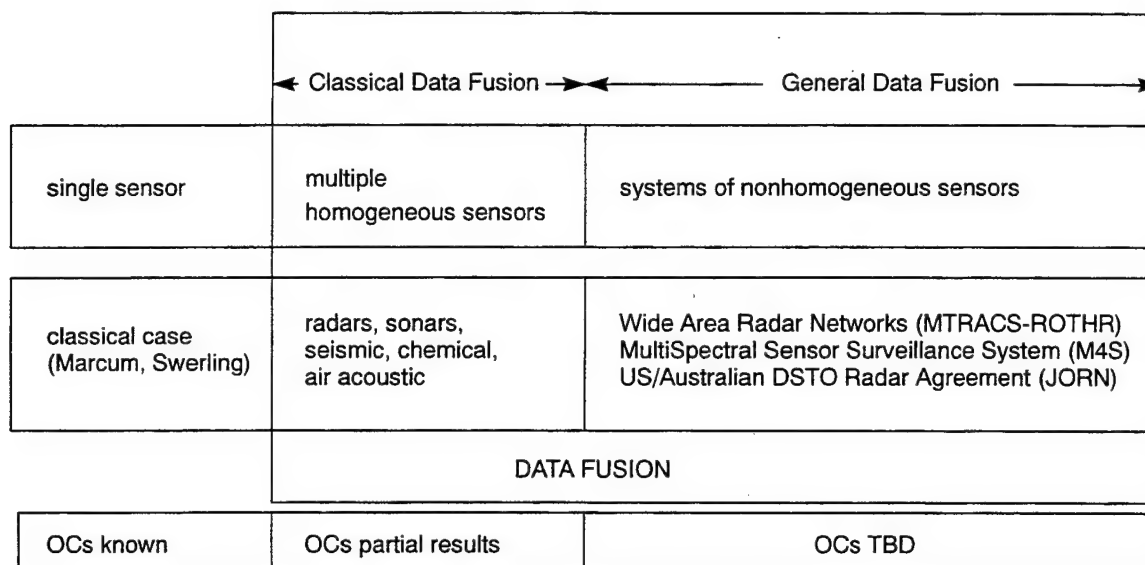


Figure 7.3.1. Current State of the Mathematical Formulation of Optimal DFAs.

REFERENCES

1. Mahler, R. 1994. "Global Integrated Data Fusion," *Proc. 7th Nat. Symp. Sensor Fusion*, vol. 1.
2. Mahler, R. 1993. "Random Sets as a Foundation for General Data Fusion," *Proc. 6th JS Data Fusion Symposium*.
3. Goodman, I. R. 1987. "A General Theory for the Fusion of Data," *Proc. 1st JS Data Fusion Symposium*.
4. Torrez, W., J. Durham, and R. Trueblood. 1993. "Performance Measures for Neural Nets Using Johnson Distributions." *Proceedings of the IEEE-ICNN Conference on Neural Networks*, vol. 1. pp. 506-509, San Francisco, CA.
5. Durham, J., W. C. Torrez, and E. W. VonColln. 1992. "Performance Analysis of the Air Defense Initiative Neural Network Processor Using Data Sets from the E1 Test." TD 2462, Naval Command, Control and Ocean Surveillance Center, RDT&E Division (now SSC San Diego), (December).
6. Torrez, W. 1991. "Arctic Signal Processors," TR 1478, Naval Ocean Systems Center (now SSC San Diego), (October).

7.4 SECTION REFERENCE

- Torrez, W. and W. Yssel. 1997. "Performance Characteristics of Data Fusion Methods with Application to Independent Surveillance Systems," *Proc. National Symp. Sensor and Data Fusion*, pp. 313-320, MIT Lincoln Lab, (April).

APPENDIX: LIST OF ACRONYMS

ADS	Automatic Dependent System
ALICE	Always Logical in Connecting Events
AOU	Area of Uncertainty
ARD	Aximuth Range Doppler
AS	Australia
ASR	Airport Surveillance Radar
ATW	Advanced Tactical Workstation
CD	Counter Drug
CERAP	Combined En-Route Radar Approach
CFAR	Constant False Alarm Rate
CIM	Current Ionosphere Model
CR	Coordinate Registration
CRT	Coordinate Registration Table
CSD	Cross Spectral Density
DF	Data Fusion
DFA	Data Fusion Algorithm
DIR	Dwell Illumination Region
DPA	Dynamic Programming Algorithm
DSTO	Defence Science and Technology Organisation
FAA	Federal Aviation Administration
FLAMOC	Fuzzy Logic and the Management of Complexity
FTAT	Fractional Track Association Time
FTHT	Fractional Track Holding Time
GPS	Global Positioning Satellite
HDR	High-Data-Rate
HFRD	High Frequency Radar Division
HTML	Hypertext Markup Language
Hz	Hertz
IOC	Initial Operational Capability
JDL	Joint Directors of Laboratories
JFAS	Jindalee Facility—Alice Springs
JORN	Jindalee Over-the-Horizon Radar Network
LOGAMP	Logarithmic Amplitude
LRID	Logical Record Identifier
MOA	Memorandum of Agreement
MOE	Measures of Effectiveness
MRTDF	Multiple ROTHRA Track Data Fusion
MTRACS	Multiple Input Tracking and Control Subsystem

NCCOSC	Naval Command, Control and Ocean Surveillance Center
NRaD	NCCOSC RDT&E Division
NRC	Nichols Research Corporation
OCC	Operational Control Centers
OTHB	Over-the-Horizon Backscatter
OTHR	Over-the-Horizon Radar
PA	Project Arrangement
PA	Propagation Assessment
PDF	Probability Density Function
PR	Puerto Rico
RDT&E	Research, Development, Test, and Evaluation
REDS	ROTHR Enhancement Demonstration System
RF	Radio Frequency
ROTHR	Radar Tracks with Relocatable Over-the-Horizon Radar
RTRT	Real-Time Raytracing
SNR	Signal-to-Noise Ratio
SPAWAR	Space and Naval Warfare Systems Command
SSC	Space and Naval Warfare Systems Center
T&E	Test and Evaluation
US/AS	United States/Australia
WARF	Wide Aperture Radar Facility
WASD	Wide Area Surveillance Division
WRF	Waveform Repetition Frequency
www	World Wide Web

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 1998		3. REPORT TYPE AND DATES COVERED Final	
4. TITLE AND SUBTITLE DATA FUSION PROJECT ARRANGEMENT: FINAL REPORT The United States/Australia Memorandum of Agreement on Radar Activities				5. FUNDING NUMBERS OMN AN: DN305088	
6. AUTHOR(S) W. C. Torrez, editor; W. J. Yssel, editor					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Space and Naval Warfare Systems Center San Diego, CA 92152-5001				8. PERFORMING ORGANIZATION REPORT NUMBER TD 3043	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Fleet Surveillance Support Command 1298 Olympic Avenue Chesapeake, VA 23322-4098				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report describes the accomplishments of the Over-the-Horizon Radar (OTHR) Data Fusion (DF) Project Arrangement (PA), initiated under the auspices of the United States/Australia (US/AS) Memorandum of Agreement (MOA) for Cooperation on Radar Activities. This report also contains Project Arrangement background information and an explicit statement and discussion of each of the three primary Project Arrangement Objectives. The three primary objectives of the Project Arrangement are as follows: 1. Share OTH radar data collected and exchange relevant technical information. 2. Jointly formulate suitable data fusion approaches, developing algorithms capable of exploiting overlapping coverage. 3. Jointly devise and conduct experiments to test and evaluate data fusion algorithms.					
14. SUBJECT TERMS Mission Area: Surveillance Over-the-Horizon Radar (OTHR) target tracking track data fusion				15. NUMBER OF PAGES 198	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED		20. LIMITATION OF ABSTRACT SAME AS REPORT

21a. NAME OF RESPONSIBLE INDIVIDUAL W. C. Torrez	21b. TELEPHONE <i>(include Area Code)</i> (619) 553-2020 e-mail: torrez@spawar.navy.mil	21c. OFFICE SYMBOL Code D7212

INITIAL DISTRIBUTION

Code D0012	Patent Counsel	(1)
Code D0271	Archive/Stock	(2)
Code D0274	Library	(2)
Code D027	M. E. Cathcart	(1)
Code D0271	D. Richter	(1)
Code D7212	W. C. Torrez	(15)

Defense Technical Information Center
Fort Belvoir, VA 22060-6218 (3)

SPAWARSYSCEN Liaison Office
Arlington, VA 22202-4804

Center for Naval Analyses
Alexandria, VA 22302-0268

Navy Acquisition, Research and Development
Information Center (NARDIC)
Arlington, VA 22244-5114

GIDEP Operations Center
Corona, CA 91718-8000